RASPBERRY Pi 5 GUIDE PG.48

STEAM DECK OLED REVIEWED PG.76

WEBCAMS RATED PG.12

MAXIMUMPC

MINIMUM BS • FEBRUARY 2024 • www.maximumpc.com

BUILD THIS PC
STEP-BY-STEP GUIDE
PG.16

MICRO BUILD, MAX POWER!

- Core i7-14700K & RTX 4070
- Mini ITX Form Factor
- Only $1,800!

DLSS Vs FSR: WHO WINS?
PG.32

HOW TO USE PASSKEYS
PG.68

FUTURE
VOL 29, NO 2

EVGA
EVGA
36

PC GAMER
EDITOR'S CHOICE
★★★★★
MAXIMUM PC
KICK ASS!

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM
Digital Storm PCs featuring Intel® Core™ i7 processors.

intel
CORE i7
inside

The Digital Storm Logo, and 'World's Most Advanced PCs' are trademarks of Digital Storm.Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

The Electronic Frontier Foundation is the leading nonprofit organization defending civil liberties in the digital world. Founded in 1990, EFF champions user privacy, free expression, and innovation through impact litigation, policy analysis, grassroots activism, and technology development. We work to ensure that rights and freedoms are enhanced and protected as our use of technology grows.

EFF.ORG

ELECTRONIC FRONTIER FOUNDATION

Protecting Rights and Promoting Freedom on the Electronic Frontier

inside

FEBRUARY 2024

SCAN TO GET THE
TOM'S HARDWARE
WEEKLY NEWSLETTER

QUICKSTART

8 **THE NEWS**
OLED monitors get upgrades; New York Times sues OpenAI; Intel's new roadmap.

11 **TECHTALK**
Jarred Walton gives the lowdown on Nvidia's new Super models.

12 **THE LIST**
The five best webcams.

Elgato's webcam is positioned as a premium camera for streamers who aren't quite ready to invest in more expensive options.

40 **SUBSCRIBE TODAY**
Subscribe to *Maximum PC* and instantly get access to over 100 back issues.

R&D

59 **HOW TO**
Browse the web with Vivaldi; Secure your secrets with VeraCrypt; Using Bitwarden.

LETTERS

14 **DOCTOR**

94 **COMMENTS**

16 **THE 14TH GEN i7 ITX BUILD**

© ELGATO, UBISOFT

32 **UPSCALING SHOWDOWN**
Jarred Walton on how DLSS, FSR, and XESS are changing things.

42 **GUIDE TO PODMAN**
Nick Peers reveals how to up your container game with Podman.

48 **RASPBERRY PI 5**
Want to spend just $60 on your next PC? *Nik Rawlinson* tells all.

54 **WINDOWS ON APPLE**
Darien Graham-Smith on how to run PC software on Apple Silicon.

IN THE LAB

74 **INTEL CORE I7-14700K**

82 **MSI MEG Z790 ACE MAX**

87 **32GB XPG LANCER BLADE RGB DDR5**

89 **WD BLACK SN770M PCIE 4.0 M.2 SSD**

90 **PRINCE OF PERSIA: THE LOST CROWN**

AMAZING

SUBSCRIPTION OFFERS

SUBSCRIBE TO ANY OF OUR **BEST-SELLING** MAGAZINES

FREE LIONS WALLCHART INSIDE
RUGBY WORLD
YOUR 2021 LIONS

MBAPPE • DE BRUYNE • RAMOS • MANCINI • MOUNT
FourFourTwo
EURO 2020
THE WAIT IS OVER

THE NO.1 MAGAZINE FOR PLAYSTATION
PLAY
DEAD SPACE
FAR CRY 6
BATTLEFIELD 2042
MOSS: BOOK II

retro GAMER
THE STAR WARS SAGA
FROM DARK FORCES TO KOTOR, WE TALK TO THE CREATORS WHO SHAPED A GALAXY FAR, FAR AWAY

TOTAL Guitar
BACK TO LIVE!
GET GIG-READY AND LEVEL UP YOUR PLAYING
GOJIRA
THE DARKNESS
MANIC STREET PREACHERS
SQUID

GOLF MONTHLY
Beat your best score
2021 RYDER CUP PREVIEW

SHOP THE FULL RANGE

WWW.MAGAZINESDIRECT.COM/**SUBSCRIBE**

Offer open to new subscribers only. Please allow up to six weeks for delivery of your first subscription issue. Payment is non-refundable after the 14 day cancellation period unless exceptional circumstances apply. For full terms and conditions, visit www.magazinesdirect.com/terms. For enquiries please call: +44 (0) 330 333 1113. Lines are open weekdays between 9am-5pm UK Time or e-mail: help@magazinesdirect.com. Calls to 0330 numbers will be charged at no more than a national landline call, and may be included in your phone provider's call bundle.

Meet the technology experts

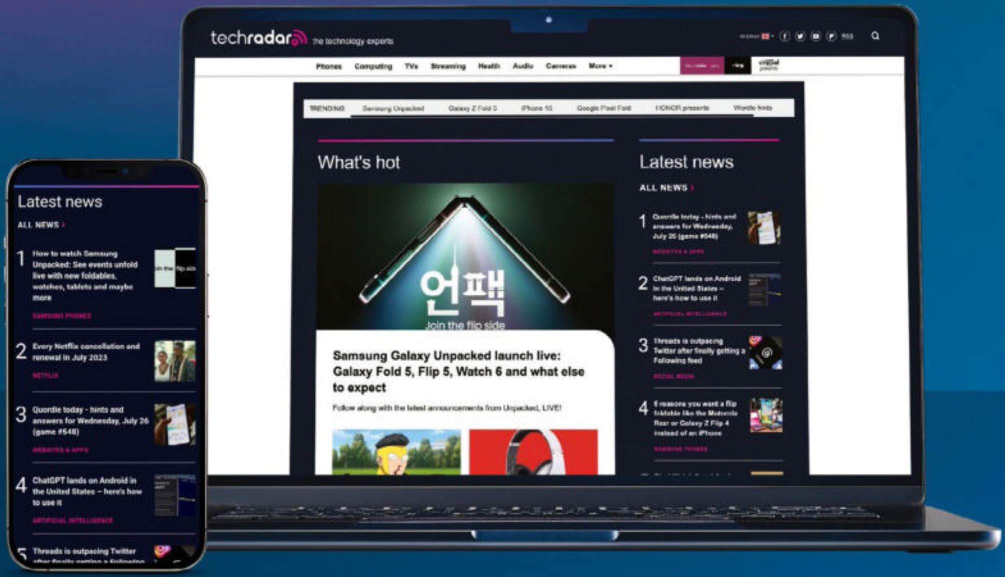
techradar the technology experts
What's hot
Latest news
Samsung Galaxy Unpacked launch live: Galaxy Fold 5, Flip 5, Watch 6 and what else to expect
Join the flip side

- The world's **most** comprehensive technology website
- An **unrivalled** mix of news, opinions, reviews and features
- **All-new design**, new homepage, new features and special reports
- Backed by **over 300 years** of editorial experience

techradar
the technology experts

www.techradar.com

MAXIMUMPC

EDITORIAL
Editor-in-Chief: Guy Cocker
Contributing Writers: Tyler Colp, Nate Drake, Ian Evenden, Darien Graham-Smith, Jeremy Laird, Chris Lloyd, Nick Peers, Zak Storey, Mollie Taylor, Jarred Walton
Production Editor: Steve Wright
Editor Emeritus: Andrew Sanchez

ART
Art Editor: Fraser McDermott
Photography: Neil Godwin, Olly Curtis, Phil Barker
Cover Photo Credits: Logitech, Future PLC

BUSINESS
US Marketing & Strategic Partnerships: Stacy Gaines, stacy.gaines@futurenet.com
US Chief Revenue Officer: Mike Peralta, mike.peralta@futurenet.com
East Coast Account Director: Brandie Rushing, brandie.rushing@futurenet.com
East Coast Account Director: Michael Plump, michael.plump@futurenet.com
East Coast Account Director: Victoria Sanders, victoria.sanders@futurenet.com
East Coast Account Director: Melissa Planty, melissa.planty@futurenet.com
East Coast Account Director: Elizabeth Fleischman, elizabeth.fleischman@futurenet.com
West Coast Account Director: Austin Park, austin.park@futurenet.com
West Coast Account Director: Jack McAuliffe, jack.mcauliffe@futurenet.com
Director, Client Services: Tracy Lam, tracy.lam@futurenet.com

MANAGEMENT
CEO: Jon Steinberg
MD Tech: Paul Newman
Group Editor-in-Chief: Graham Barlow
Group Art Director: Warren Brown

PRODUCTION
Head of Production: Mark Constance
Senior Production Manager: Matthew Eglinton
Production Manager: Vivienne Calvert
Production Assistant: Emily Wood

Future US LLC, 130 West 42nd Street,
7th Floor, New York. NY 10036. USA. www.futureus.com

INTERNATIONAL LICENSING & SYNDICATION
Maximum PC is available for licensing and syndication.
To find out more, contact us at licensing@futurenet.com or view our available content at www.futurecontenthub.com.
Head of Print Licensing: Rachel Shaw, licensing@futurenet.com

SUBSCRIBER CUSTOMER SERVICE
Website: www.magazinesdirect.com
Tel: 844-779-2822
New Orders: help@magazinesdirect.com
Customer Service: help@mymagazine.co.uk

BACK ISSUES
Website: https://bit.ly/mpcsingleissue

Next Issue On Sale February 27, 2024

EXTRA DIGITAL FEATURES
AUDIO FILE
PHOTO GALLERY
VIDEO FILE

© 2024 Future US, Inc. All rights reserved. No part of this magazine may be used or reproduced without the written permission of Future US, Inc. (owner). All information provided is, as far as Future (owner) is aware, based on information correct at the time of press. Readers are advised to contact manufacturers and retailers directly with regard to products/services referred to in this magazine. We welcome reader submissions, but cannot promise that they will be published or returned to you. By submitting materials to us, you agree to give Future the royalty-free, perpetual, non-exclusive right to publish and reuse your submission in any form, in any and all media, and to use your name and other information in connection with the submission.

FUTURE
Connectors.
Creators.
Experience
Makers.

Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR)
www.futureplc.com

Chief Executive Officer **Jon Steinberg**
Non-Executive Chairman **Richard Huntingford**
Chief Financial and Strategy Officer **Penny Ladkin-Brand**

Tel +44 (0)1225 442 244

© GETTY IMAGES

Guy
Cocker

SUPER GPUs

AS I WRITE THIS, I've just returned from this year's Consumer Electronics Show in Las Vegas, where I saw enough laptops, desktops, and monitors to easily fill these pages for the next 12 issues. Coming away knowing a bit more about the lay of the land in PC tech for 2024, I can say that the future looks exciting, even if there's no identifiable 'next big thing' set to revolutionize the industry this year.

The clear highlight for most will be the release of Nvidia's SUPER series of graphics processors. The RTX 4080, RTX 4070 Ti, and RTX 4070 SUPER cards don't look to be huge upgrades in terms of performance, but they do look to right some of the wrongs of the vanilla 40-series cards when it comes to pricing, particularly the RTX 4080, which comes in at $999. That's $200 less than its predecessor, and really takes the fight to AMD's RX 7900 XTX, which we've already seen discounted by retailers to $800.

Intel's 14th generation CPUs also made their way to laptops, OLED panels were being pumped out by every monitor manufacturer going, and the MSI Claw showed that Intel was ready to get into handheld PC gaming. There are lots of mentions of CES in this issue, but I'm going to be putting my Best of CES report containing all the show's top PC tech for next issue, which will be on sale February 27.

As for this issue, our master builder Zak Storey has put together a micro PC for the first time in over three years, when his last one adorned our October 2020 cover. Back then, the brief was a 10th-gen Intel CPU and an Nvidia RTX 20-series GPU mounted to a Hydra Case. This time, it's a 14th gen Intel and a 40-series Nvidia going into (well, onto) the same chassis. It's our first build with the Core i7-14700K (which you can also find heading up this month's reviews on page 74), with the brief being: can you air cool Intel's chips, given they touch 100 C under load? And can you build a PC that can literally be thrown in a backpack, given how large the latest Nvidia GPUs are? All these questions are answered in Zak's cover feature, starting on page 16.

Elsewhere, we review the Steam Deck OLED (page 76) and find the portable PC gaming market in rude health two years since Valve kickstarted the whole thing with its debut model. We also take a look at the new Raspberry Pi 5 on page 48, and find a machine that can do more than just automate your smart home gear or play emulators like its predecessors. This machine is powerful enough to be a bonafide desktop replacement for those who don't mind tinkering with Linux.

This month's issue also holds one of Jarred Walton's deep dive features, where on page 32 he looks at the current state of upscaling tech from the big three: Nvidia, AMD, and Intel. With some of these technologies now in their third generation, it was time to take a look at how they're faring, and which one you should be using. As usual, Jarred has gone in-depth in his comparisons, so if you love a well-researched graphics card feature as much as I do, I think you'll enjoy it.

We also have a wealth of tutorials, from passkeys in Bitwarden, to running Windows on a Mac, and hiding your OS with Veracrypt.

Enjoy the issue!

Guy Cocker

Guy is Maximum PC's *editor-in-chief. He built his first gaming PC in 1997 to play* Tomb Raider *on 3dfx, and has been obsessed with all things PC ever since.*

↘ submit your questions to: editor@maximumpc.com

THE NEWS

Nvidia goes Super

We've been expecting you

NVIDIA'S 40-series card brought us the technical marvel that is the RTX 4090. Sure, it chewed through power, occasionally melted cables, and cost $1,599. It hogged the limelight, and we all wanted one. Further down the 40-series range, things weren't as rosy. The RTX 4080 was deemed overpriced. Anybody who was after the ultimate gaming card sprang for a 4090 if they could. Others looked around for a better deal, or waited.

The RTX 4070 Ti was originally going to be a 12GB 4080, but was 'unlaunched', as Nvidia claimed, because the performance didn't warrant the '80' brand. As a 4070 Ti, it was thought pricey compared to the previous generation's 3070 Ti. Further down the range, Nvidia faced stiff competition from AMD, and the 40-series struggled to be competitive. Nvidia was expected to fix this—now, it finally has. We are to get three new cards, all branded 'Super'. This goes a long way to addressing the range's haphazard scaling. If you've just bought one of the original cards, you have every right to feel annoyed—these are either faster, cheaper, or both. At least Nvidia made sure people knew they were coming.

RTX 4080 SUPER

It's hard to complain when something's faster or cheaper. Nvidia's Super series fixes the 40-series' quirky line-up.

The RTX 4080 is to be replaced by the RTX 4080 Super. This will cost $999, which is a significant price drop. We also get a small GPU upgrade—the full-die AD103 chip has 10,240 CUDA cores rather than 9,728. The memory is also faster: 2.7 percent at 23GB/s. These changes will only translate in a few frames at best—no doubt Nvidia felt obliged to increase the specification to earn the 'Super' badge. Basically, this is a RTX 4080, but $200 cheaper.

Below that sits the RTX 4070 Ti Super, which replaces the 4070 Ti. Here, we get more CUDA cores—8,448 instead of 7,680. We also get 16GB instead of 12GB, which now sits on a 256-bit bus rather than a 192-bit one thanks to an upgrade from the AD104 chip to the AD103 used in the 4080 cards. This is a decent bump in hardware, and enough performance to take a stab at 4K. Nvidia claims it is 2.5 times faster than an RTX 3070 Ti, albeit with frame gen enabled. All this for $799.

Lastly, we have the RTX 4070 Super. Its AD104 GPU now has 7,168 CUDA cores—up from 5,888—and over 20 percent more cores. This means 36 TFLOPS against 29, and it's claimed to be faster than an old RTX 3090. It also means more power, going from 200W to 220W. The memory stays the same—12GB on a 192-bit bus. Unlike the two other Super cards, this doesn't replace the original, with the RTX 4070 cut to $549.

Nvidia also has a slower version of the RTX 4090, the 4090D. This has been made for the Chinese market. The standard RTX 4090 was banned for export there last October, and this is Nvidia's response. It draws 425W rather than 450W, and has 1,792 fewer CUDA cores—about five per cent slower. It has been designed to be compliant with export restrictions.

Nvidia has come in for some heat for working around sanctions like this—it was publicly warned by the US Commerce Secretary that any cards remade to pass sanctions will be subject to scrutiny, and possible banning. Nvidia has obeyed the letter, if not the intent of the law. Not that this worries Nvidia—it's making a fortune selling high-end cards and AI accelerators. It reportedly has 90 percent of China's AI chip market, worth over $6 billion a year.

Meanwhile, at AMD, we have the new Radeon RX 7600 XT. This has 16GB, but still on a 128-bit bus running up to 288GB/s. The GPU is the same Navi 33 chip with 32 RDNA 3 Compute Units, which means 2,048 stream processors. Clock speeds are up 220MHz to 2.47GHz.

That extra memory means you should be able to set everything to maximum for 1080p gaming, and maybe play at 1440p. That's the idea, anyway—the memory bus is a bit slow to fully service that much memory, and Navi 33 is really aimed at 1080p. Whatever, the RX 7600 XT costs $327—a $60 hit over the vanilla version. **–CL**

> The three new cards go a long way to addressing the range's haphazard scaling

© NVIDIA

OLD MEDIA SUES AI

THE NEW YORK TIMES has become the first big media organization to sue OpenAI over copyright infringement. The lawsuit claims OpenAI has been trained on content it produced, and now competes with them using the results. It claims OpenAI should be liable for billions in statutory and actual damages, giving examples of ChatGPT quoting and paraphrasing copyright material.

The ire of the *New York Times* is understandable, but can this go anywhere? Copyright laws are lagging behind the rapid development of AI, and there's no clear precedence. The whole AI industry has been lax over where it gets training material. Haphazardly scraping it off the internet means masses of copyright material has been fed into the machine, so we shouldn't be surprised that the output looks suspiciously familiar. Purging the models of such material would be impossible—besides, it needs masses of recent data to be useful. Some sort of compromise will have to be reached. But right or wrong, suing a company with pockets as deep as Microsoft has is a tough job. **–CL**

GOOGLE KILLS TRACKERS

Google
chrome

ABOUT ONE PER CENT of Chrome users—that's 30 million people—have had third-party cookies disabled, dubbed Tracking Protection. This is part of a project by Google that will have major implications on internet advertising. Google will slowly increase that number over the next few months, and by the end of the year it'll be everybody.

Third-party cookies are beloved by advertisers, but not always by the rest of us, as our activities are monitored so we can be targeted more effectively. Google has been under pressure to do something for a while, and announced in 2020 that it would, even promising not to invest in any alternative tracking systems. Now, it is putting its plan into action. You will have the option to turn them back on—some sites will get obstreperous if you don't. Cookies from individual websites are okay—it's those third-party trackers that are being culled. The advertising industry will just have to adapt, and it will. **–CL**

Tech Triumphs and Tragedies

A monthly snapshot of what's good and bad in tech

TRIUMPHS

✓ **COPILOT GETS OWN KEY**
Microsoft wants AI on the laptop keyboard via a dedicated Copilot key—the first change since 1994's Start key.

✓ **GAME OVER**
NES classic *Tetris*'s kill screen is somewhere above level 155. One player hit it at level 157.

✓ **WI-FI 7 CERTIFIED**
The Wi-Fi Alliance is ready to certify Wi-Fi 7 devices. 320MHz enables 46Gbps, nearly five times the speed of Wi-Fi 6.

TRAGEDIES

✗ **HYPERLOOP IS DEAD**
Mr Musk's idea to build a high-speed train using magnetic levitation was interesting, but unworkable in practice.

✗ **VR SALES COLLAPSE**
Sales of VR hardware are down 40 percent. Meta's losses in VR over the last two years top $25 billion.

✗ **BAD YEAR TO BE A TECHIE**
Research claims only 700 IT jobs were added in 2023, against 267,000 in 2022.

© GOOGLE, WIKIMEDIA, MSI

MSI CHOOSES INTEL TO RIVAL STEAM DECK

Behold the new Claw

MSI BUILDS LAPTOPS and now PC gaming handhelds too. The new Claw promises smooth 1080p gaming on AAA titles, thanks to the first use of an Intel Core Ultra processor in such a device. At its heart is an Intel Core Ultra 7 155H, a Meteor Lake hybrid design. It boasts six performance cores and eight efficiency cores. The P cores max out at 4.8GHz, the E cores at 3.8GHz. It has integrated Arc graphics, with eight Xe cores, and XeSS upscaling. Intel's recommended customer price is $503. MSI isn't the first to try Intel in a gaming handheld, but it is the first big player to go Team Blue for this round.

You view the action on a 7-inch touchscreen (1,920 by 1,080) running at 120Hz. There's a PCIe 4.0 M.2 slot for storage, 16GB of memory, and a Thunderbolt 4 port. The 53Wh battery is claimed to be good for two hours of serious gaming. There's plenty of cooling on offer, with two fans, two heat pipes, and lots of vents. The Intel chip has an official TDP of 28W, but can go much higher, with the full boost quoted as 115W by Intel. It follows the basic design of handhelds: lots of buttons, twin analog thumbsticks, and a weight of 675g. A dock with lots of ports and an external GPU will follow.

The rivals are Valve's Steam Deck, but there's also Lenovo's Legion Go, and Asus' ROG Ally. All use AMD silicon—MSI reckons the Claw's Intel innards give it the edge. The $64,000 question is whether it will. Speaking of money, the Claw is due this spring, and there will be three versions. The $699 model has 512GB storage, and a Core Ultra 5 135H processor with four fewer performance cores. Then there's the $749 model with the full Ultra 7, and the $799 model, with 1TB of storage. Expensive, but it looks sweet. **–CL**

MSI's new gaming handheld offers powerful Intel silicon, swish design, and a big battery.

Odyssey OLED G8

OLED MONITORS GET UPGRADES

You know you want one

WHILE THEY LOOK LOVELY, OLED screens have been something of a luxury. But prices have been falling, and we've reached the third generation of panels, with numerous releases timed for CES. They're bigger and brighter, but burn-in still exists, even if it's nowhere near the problem it was once feared to be.

Just before the holidays, LG announced five new UltraGear models, including the 32GS95UE, a 32-inch screen that can switch between 4K at 240Hz, and 1080p at 480Hz. Now, Samsung has announced the Odyssey G6 G60SD, G8 G80SD, and the G9 G95SD, the latter more of an update than an all-new panel. The smallest is the 27-inch G6, a 2,560 by 1,440 panel with a game-friendly refresh rate of 360Hz, and a response time of 0.03ms. The G8 is a 32-inch 3,840 by 2,160 screen (4K), with a refresh of 240Hz, and the same 0.03ms response. The G9 is a curved 49-incher with a 32:9 aspect ratio at 5,120 by 1,440.

All support VESA DisplayHDR True Black 400, a certification optimized for OLED panels that takes into account the nuances of panel brightness over the whole screen. That G8 looks to be the sweet spot—4K means a sharp pixel density of 140DPI. Others joining the third-gen OLED party include the Alienware 32 AW3225QF, another 32-inch 240Hz 4K panel, and likely to be the first you can buy.

Asus ROG has a new trio too, including the dual-mode ROG Swift PG32CDP that also switches between 240Hz 4K, and 480Hz 1080p (LG's panel). There's also a 39-inch curved ultrawide PG39WCDM, and the PG27AQDP, the world's first native 480Hz OLED aimed at 1440p (another LG panel). None of these monitors will be cheap, but with the wide range of models available, here's hoping the competition helps to bring prices down. **–CL**

Intel's process push

INTEL'S CEO, Pat Gelsinger, is an engineer, and when he got the job back in 2021, he had an engineer's solution to the company's floundering development schedule. He released a roadmap that promised five nodes in four years. The plan hasn't gone entirely smoothly: that Raptor Lake refresh wasn't part of the deal, Intel 4 barely made it on time, and both the 20A and 18A 'Angstrom era' nodes seem in the balance.

In a few weeks, we'll have another roadmap. Intel plans to get a trillion transistors on a 'package' by 2030. The company still plans to jump to the cutting-edge in barely four years. This year, we'll allegedly see Intel 3 in volume, with 20A production supposedly by the end of the year. Part of the strategy has always been to build chips for others. TSMC is about to move into serious 3nm volumes—the target is apparently 80 per cent of all production this year to be its second generation N3E chips. Apple pretty much bought the first generation for its A17 Pro and M3 processors.

Anyway, TSMC's roadmap has been updated to 2030. It shows 2nm in 2025, 14A by around 2028, followed by 10A in 2030. It also talks of trillions of transistors on stacked packages, and up to 200 billion on a chip. Intel's plans come courtesy of ASML's extreme ultraviolet lithography machine, the EUV 0.55 NA, which is what you need to really move into Angstroms rather than nanometers. Pundits don't expect TSMC to switch for a while yet. If Gelsinger's promises looked like hubris, the next 12 months will be the real test. If Intel can hit its targets, it could be the most advanced chip producer in the world. But failure could be existential for the firm. **–CL**

Intel has cooked up some seriously ambitious plans.

New Windows, same files

A Windows installation can get bloated. Wouldn't you like a fresh OS without losing your stuff? Microsoft's latest beta has an option dubbed 'Fix Problems Using Windows Update'.

It installs a new copy of the version you are running without removing your applications or files. You may need to run updates you've missed, which arguably is not new. Previously, the process required you to make a bootable flash drive, then be careful about what you selected. This option is one-button. It's a candidate for inclusion in the next update. Microsoft has been quiet about this, probably because it's the sort of thing that has definitely got to work first time, and maybe that isn't quite the case just yet. **–CL**

Apple's VR gamble

Apple has confirmed a launch date of February 2 for the Vision Pro headset. For $3,499, you get an M2 processor, 4K OLED displays, and plenty of cameras. It's a unique blend of augmented reality and virtual reality. Apple arrived late to the VR party, but has something characteristically innovative.

Software-wise, visionOS has a three-dimensional interface operated by voice or movement. There are immersive environments to gawp at, and it even shows your face on an external screen. You don't get much battery life, though—about two hours tether-free. Apple's system is stunning, but is hardly mainstream at this price. Has it made something we didn't know we really wanted again? **–CL**

© SAMSUNG, INTEL

Jarred Walton

TECH TALK

Nvidia RTX 40-series Super models take the runway

NVIDIA'S RTX 40-SERIES Ada Lovelace GPUs have had a rocky reception, mostly on account of their significantly higher generational prices, reduced memory interface widths, and reliance on DLSS 3 frame generation to make performance look better.

It's all reminiscent of the RTX 20-series, which had similar complaints. Just like the 20-series, Nvidia will trot out a mid-cycle refresh in the form of new Super models.

The RTX 4070 Super has specs that put it close to the existing RTX 4070 Ti. It will feature 56 Streaming Multiprocessors (SMs) and 7,168 CUDA cores—just seven percent below the 4070 Ti and 22 percent above the RTX 4070. More importantly, it will inherit the $599 MSRP of the RTX 4070, which will get a price cut to $549. The memory will be the same 12GB of 21 Gbps GDDR6X, with a 220W TGP (total graphics power) rating—20W more than the 2070.

RTX 4070 Ti Super comes next, combining the Ti and Super suffixes to offer a new experience for Nvidia GPU names. The good news is that the 4070 Ti Super will fix one of the major issues we had with the RTX 4070 Ti, as it will leverage the AD103 GPU and offer a 256-bit memory interface and 16GB of VRAM, with 33 percent more bandwidth. Otherwise, it's only a minor upgrade, with 10 percent more shader cores. It will take over the same $799 price point as the RTX 4070 Ti, which will get phased out. Like the 4070 Ti, the Super variant will only be offered from Nvidia's add-in card partners—no Founders Edition. The base TGP remains 285W.

RTX 4080 SUPER
RTX 4070 SUPER

Nvidia is set to launch three new 'Super' models in 2024: the RTX 4070 Super, RTX 4070 Ti Super, and RTX 4080 Super.

© NVIDIA

> All three cards represent modest improvements in price and performance

Lastly, the RTX 4080 Super will use the same AD103 GPU as both the vanilla 4080 and the 4070 Ti Super, and still be limited to a 256-bit memory interface and 16GB of GDDR6X memory. It gets a bump in shader counts from 76 SMs in the 4080 to the full 80 SMs in the 4080 Super. It also has memory clocked at 23 Gbps, another boost that will be phased out once the Super variant arrives. Given the lack of massive changes, the TGP remains at 320W as well. The RTX 4080 Super will have a $999 MSRP—$200 less than the outgoing card.

I can't help but feel disappointed with the 4080 Super, as I hoped to see a cut-down AD102 implementation and 20GB of memory. Still, all three cards represent modest improvements in price and performance. We've already seen several of Nvidia's RTX 40-series drop below their official MSRPs, and the new Super variants may push prices lower.

The RTX 4070 still needs to go lower, as the 4070 Super will theoretically offer about 20 percent higher performance for nine percent more. When I spoke to Nvidia, it said that retail prices don't have to follow the MSRPs. Nvidia doesn't want to officially cut the price that much, but we could see $499 RTX 4070 cards soon.

Overall, the 40-series Super announcement delivers what you'd expect from a mid-cycle refresh. Anyone who already owns a 40-series GPU will be fine, but if you were on the fence about upgrading before, the new models may entice you. At the same time, we're one year closer to the next-generation Blackwell RTX 50-series GPUs. Consumer models will likely arrive in 2025, though whether Blackwell will deliver major improvements in performance and value remains an open question.

Jarred Walton has been a PC and gaming enthusiast for over 30 years.

THE LIST

THE BEST WEBCAMS

WHAT EXACTLY do you want from a webcam? How about something that makes you look your best, no matter how dimly lit your bedroom or home office? Just be aware that there are webcam horses for specific streaming courses. The best webcam for live streaming isn't necessarily the same as the best option for remote working. Many of these webcams support HD and 4K up to 60fps, for instance, but the latter in particular adds to the price tag, and might be overkill for daily office duties.

5 DELL ULTRASHARP WB7022

This is one of the best 4K webcams you can buy. On top of offering 4K at 30fps, you've got HDR support and even AI-powered auto framing. As expected, UltraSharp's support for 4K recording delivers impressive detail over your standard 1080p webcam. It works well in poorly lit or overexposed rooms—just note that if you're only using this for work calls, the webcam's picture quality might be nerfed by apps like Meet or even Zoom and their aggressive video compression. **$185, www.dell.com**

4 RAZER KIYO PRO

The Kiyo Pro improves on Razer's previous webcams, with a slightly different approach to handling lighting in dark spaces. It ditches the ring light in favor of a sophisticated light sensor to tackle the gloom. The Kiyo Pro's other improvements include HDR (off by default), 1080p resolution at 60fps, a wide-angle lens, and omnidirectional microphone. It's a feature-packed webcam for streamers. Night-time or darker rooms is where the Kiyo Pro truly shines—it's one of the best low-light webcams we've ever used. **$80, www.razer.com**

3 LOGITECH STREAMCAM

Designed for—you guessed it—streamers and other content creators, you can rotate this on its three-axis clip for portrait mode. The mount can also be angled face-down up to 90 degrees for keyboard and mouse movements, while an extra mount allows screws for more complex arrangements. It shoots in 1080p at 60fps, and the Capture 2.0 software automates most settings. **$95, www.logitech.com**

2 LOGITECH C922 HD PRO

Unless you want specific features, there is no better value than this. Its sharp 1080p images, paired with a wide field of view and great autofocus, make it a fantastic video-conferencing choice. Low-light performance is great; the noise levels don't shoot through the roof if you turn off a few lights. Most of the settings can be adjusted through Logitech's Camera app, and streamers will appreciate its compatibility with ChromaCam background replacement. **$69, www.logitech.com**

1 ELGATO FACECAM

Elgato's webcam is positioned as a premium camera for streamers who aren't quite ready to invest in more expensive options, such as a DSLR camera, but are still looking for the best picture they can get. The Facecam offers uncompressed video at 1080p/60fps, which is a pretty huge deal and means you have final video output with less artifacting. Out of the box, the picture quality of the Facecam is great and the latencies are low, while the Camera Hub software lets you adjust your camera settings with ease. **$199, www.elegato.com**

© DELL, RAZER, LOGITECH, ELGATO

Jeremy Laird

TRADE CHAT

Is Intel's roadmap for real?

BY THE TIME you read this, Intel will have rolled out a whole new roadmap, the next waypoint on its journey to renewed technological dominance. Except I can't help but notice that the company hasn't come close to delivering on its existing roadmap. You know, the one that has Intel delivering five new chip production nodes in four years?

It was February 2021 that Pat Gelsinger returned triumphantly to the helm of Intel. It was a winning narrative. Intel had lost its way after years of marketing-biased leadership. At last, an engineer was returning to the top job—and not just any engineer, but a former Intel devotee who'd joined the company aged just 18 in 1979 and worked his way up through the firm to become CTO, only to be forced out by those evil marketing suits in 2009. Look what happened to Intel since.

I was pretty down with Gelsinger's return, too. His keynotes were a definite highlight back in the good old days of pure print journalism and trips to San Francisco for the Intel Developer Forum. Well, Gelsinger has now been in the top job for three years, and we are rapidly approaching judgement day. Of course, the whole 'five nodes in four years' thing was a bit specious to begin with. Those five new nodes included Intel 7, Intel 4, Intel 3, Intel 20A, and Intel 18A. At best, that list only contains three truly new nodes, arguably only covering two.

Intel 7, of course, is a rebranded version of the 10nm node that predated Gelsinger's return to the company, so that's not new. Intel 4, therefore, is the first genuinely new node, with Intel 3 a revision thereof. Intel 20A is the next actual new node, and 18A is, again, a tweak of 20A. So that's Intel 4 and Intel 20A as the two new nodes, thus what initially seemed like an insane development roadmap turns out to be standard Moore's Law fare involving a die shrink every two years.

The problem is that it's debatable if Intel is going to manage that. The company just about managed to keep things on schedule and get its new Meteor Lake CPUs out the door in 2023. But when you take a close look at Meteor Lake, you'll find that only one chiplet out of five inside the package is built on Intel 4. Meanwhile, Intel had to refresh Raptor Lake on the desktop using old Intel 7 silicon, which surely wasn't part of the original plan.

> “I'm not saying it won't happen, but it definitely doesn't look promising right now

Intel's Meteor Lake laptop chips have barely gone on sale, but will supposedly be replaced soon by Lunar Lake.

© INTEL

To keep on track, Intel needs to sell chips produced on not only the revised Intel 3 node, but also the new Intel 20A node. I'm not saying it won't happen, but it definitely doesn't look promising right now. Intel put on a brave face at CES, reaffirming its commitment to ship the next-gen 20A Arrow Lake desktop CPU architecture later this year, though probably only the CPU tile within the package will be on 20A.

Intel also said that Lunar Lake for laptops will arrive this year, again on 20A. So, the first Meteor Lake laptops are barely on sale, and already Lunar Lake chips are arriving this year? That scheduling looks so rushed, of course, because Meteor Lake arrived very late. It seems clear that Intel 4 is proving problematic, but somehow Intel 20A is going to arrive on schedule, and itself be revised and improved in time to allow for 18A chips to go on sale in 2025? I find that tricky to compute.

Admittedly, Intel does seem to be pressing on. Just days ago, Intel announced that it is the first customer to take delivery of ASML's latest and greatest Twinscan EXE:5000 High-NA EUV lithography scanner, the machine it will use to move beyond 18A and enable whatever nodes it plans to announce. But it needs to deliver on that original 2021 roadmap first, and for now, that seems to be hanging very much in the balance.

Six raw 4K panels for breakfast, laced with extract of x86... Jeremy Laird eats and breathes PC technology.

DOCTOR

THIS MONTH THE DOCTOR TACKLES…

> N100 conundrum
> Fix corrupt Office
> Turn back time

Which N100 board?

I'm looking to upgrade my J4125-ITX NAS server, and have settled on Intel's newest chip—the N100 processor, which benchmarks tell me represents a huge leap forward over the previous board, even though the J4125 chip has served me well. I see that in addition to ASRock's N100-based board (the N100DC-ITX), Asus has released its own board, the Prime N100I-D D4.

I can see two obvious differences between the boards: the N100DC-ITX has two onboard SATA connectors to the N100I-D D4's one, but the Asus board works with a standard power supply like the one currently in my J4125-ITX build, while ASRock's board has its own built-in power supply with 19V jack, into which I would need to plug in a compatible laptop adapter. Which would you recommend?

—**Michael James**

THE DOCTOR RESPONDS: Yes, the N100 is a stellar chip, delivering the equivalent performance of a 65-watt i5-7400 or i3-9100 in a 6-watt chip. Sadly, as you've identified, there are trade-offs in terms of available ports, as well as power supply choices. Ultimately, your choice may boil down to availability, as the Asus Prime N100I-D D4 is almost impossible to source from the US—we were able to track down a couple of listings on eBay that shipped from Europe (Germany and Italy respectively) for around $180, but at time of writing there was no inventory from any US-based sellers. On the other hand, the ASRock N100DC-ITX can be found for $129.99 on newegg (www.newegg.com/p/N82E16813162133).

Both boards are disappointingly bare when it comes to storage—in addition to the SATA port(s) on offer, both include a single M.2 2280 slot (PCIe 3.0x2) for running your boot drive from, so you'll need to budget an additional $50+ for an M.2 NVMe drive, such as the NAS-friendly WD Red SN700 500GB drive ($59.99, www.walmart.com/ip/154644440).

Although the N100DC-ITX comes with an extra SATA port, and there's no problem (in theory) adding two more via the M.2 Key E slot (search eBay for 'M.2 key A+E male SATA' for a suitable adapter for $10-15), ASRock discourages its usage with more than two mechanical (i.e., HDD as opposed to SSD) hard disks. This is down to the N100DC-ITX's built-in PSU, which requires connecting to a 19V laptop power adapter, and the fact you need to factor in the maximum amount of power each HDD needs when starting up—a theoretical maximum of 21W per drive.

Asus produces our choice of N100 motherboard.

© ASUS

We say theoretical, because in practice this rarely happens, even during startup. We monitored our own J5040-ITX rig with four WD Red HDDs and one SSD, and it maxed out at 50W during startup before eventually dropping to around 30W. Nevertheless, your power adapter must be able to handle the initial demand, which is why ASRock recommends 90-watt 19v adapters to provide plenty of headroom.

However, there's one more hurdle: because there's no PSU involved, power is delivered to all SATA drives from a four-pin SATA_PWR1 port on the motherboard. You will need to source a splitter to connect to it, effectively daisy chaining all your SATA drives from a single power port. The Doc would hesitate to recommend the N100DC-ITX for anyone wanting to connect more than two SATA drives, in addition to the NVMe boot drive.

Although difficult to source, we'd recommend seeking out an Asus Prime N100I-D D4 for around $160-180 including international shipping, and while you only have one onboard SATA port, its standard 12V power connector means powering additional drives is simple using your existing PSU. You could add an M.2 Key E adapter as with the ASRock board, but that would only give you three SATA ports, so you're better off adding more SATA ports via the PCIe 3.0x1 slot.

↘ submit your questions to: doctor@maximumpc.com

The PCIe SATA adapter card you choose depends on how many ports you need, which in turn informs the chipset you require. For example, if you need a six-port adapter, make sure it's based on the ASMedia ASM1166 chip, while five-port adapters should sport the JMicron JMB585 chip. A four-port adapter is usually sufficient for most, in which case make sure it's using one of ASMedia's ASM1064 or ASM1164 chips. An ASM1064-based four-port card can be had direct from newegg for just $27.57 (www.newegg.com/p/1B3-002S-00003), for example.

One final reason to choose the Prime N100I-D D4 board: because it utilizes a standard PSU, you'll save money by reusing the one in your setup, leaving you the task of sourcing suitable RAM. The Asus board sports a single SO-DIMM slot that supports DDR4-3200 memory with XMP, but while Asus claims it only supports up to 16GB RAM, the board does unofficially support 32GB. However, if you wish to stick to Asus's recommendations, there's little benefit to paying extra for XMP, so consider 16GB of Kingston's ValueRAM ($41.45, www.newegg.com/p/1X5-0009-00940), which would bring the outlay of your upgrade to just under $300 for mobo, RAM, PCIe SATA adapter, and 500GB NVMe boot drive.

Office repair

Out of nowhere, I'm unable to use any Office applications. Whichever one I launch, I get the following error: 'We're sorry, but Word (or Excel, or whatever) has run into an error that is preventing it from working correctly. Word/Excel/etc will need to be closed as a result.' It also offers to repair the program, but clicking Repair Now does nothing—occasionally I'm able to run the app for a few minutes, but then it closes without warning, and unless I've saved in the meantime, all my edits are lost.

I've tried going into System > Apps > Installed Apps, but any attempt to repair or reset the installation has no effect. I've even tried uninstalling and removing, but it still refuses to work.

—Aaron B Dick

THE DOCTOR RESPONDS: This is probably linked to a recent update where Office 365 has now been renamed Microsoft 365. In some rare cases, it ends up corrupting your installation beyond its ability to recover or repair itself. You need to source the Support Recovery Assistant tool, which can be used to remove Office from your PC before—after a reboot—offering to download and reinstall Office from scratch.

The simplest way to obtain the tool you need is to type https://aka.ms/SaRA-OfficeUninstall-sarahome directly into your browser—you'll be prompted to download SetupProd_OffScrub.exe. It's a tiny 192KB download—once done, double-click it and follow the prompts. It'll download the full tool (click Run when prompted), then launch the Microsoft Support and Recovery Assistant.

Your Office installation should be detected—check this and click Next. Office will be—possibly slowly—removed from your PC, then you'll be prompted to reboot. After Windows restarts, the assistant will offer to reinstall Office automatically for those with Microsoft 365 subscriptions (it'll offer links to those who wish to download setup files for all versions of Office from 2010). Click Yes, verify 'English (United States)' and '64-bit' are pre-selected, and check 'I have saved my work…' before clicking Install.

Office will again be downloaded and installed—this can take 30 minutes or longer, so be prepared to wait. Once complete, you should find Office working again, with most of your settings hopefully intact.

Roll back my PC

I regularly install software for testing purposes that I quickly realize I don't want, and uninstall it. I use Bulk Crap Uninstaller, which helps get rid of most traces, but I'd like some kind of snapshot system—like that offered on my VirtualBox virtual machines—where I can magically restore my PC to a specific point in time. Can you offer any suggestions?

—Christopher Winkler

THE DOCTOR RESPONDS: One solution would be to investigate the Rollback RX series of products. If you're willing to spend over $80, try the trial version of RollBackRx (https://horizondatasys.com/rollback-rx-time-machine), which would fit your needs perfectly. It can be configured to take snapshots manually or on a schedule, and restoring your system to any previously taken snapshot takes no longer than a single reboot.

There's also free entry-level tool, Reboot Restore Rx (https://horizondatasys.com/reboot-restore-rx-freeware), which allows you to undo system changes by rebooting or hard resetting your PC—on every boot, a saved baseline is restored, wiping out any changes in the meantime. But while it's possible to manually update the baseline, forgetting to do so would undo any changes made to saved documents and files. It might prove more trouble than its worth.

One solution would be to employ a drive imaging tool like Macrium Reflect Home ($79.99, www.macrium.com/products/home) or Hasleo Backup Suite Free (free, www.hasleo.com). Both support delta restore (Hasleo added this in its recent v4.0 update), which reduces image restoration times by only restoring the differences between the backup and current system.

If you were to set up a daily incremental backup plan with either, you'd give yourself a fallback where you'd only lose the last 24 hours' worth of changes if you wanted to roll your system back. However, both allow you to take more frequent incremental updates (Hasleo users can select 'Run at intervals' to take hourly backups), or you can try to get into the habit of manually taking an incremental backup prior to installing software on your PC, so rolling back would create the minimum amount of collateral damage.

Update Jellyfin

I recently made note of Jellyfin's blog post noting two vulnerabilities and recommending we update to 10.8.13 as soon as possible. However, despite religiously checking my Ubuntu Server 22.04 LTS system for updates, no release is forthcoming. Why is it taking so long to appear on Ubuntu?

—Henry Krupp

THE DOCTOR RESPONDS: It sounds like you installed Jellyfin through the default Ubuntu repositories. These are frozen at the point of Ubuntu 22.04's release, so by now you'll be several versions behind the latest available version. The fix is simple—open a Terminal window and issue the following command:

```
curl https://repo.jellyfin.org/install-debuntu.sh | sudo bash
```

This will download a script to add the official Jellyfin repository, which will replace your outdated version with the latest available build. It will continue to offer you the latest release as soon as it's available, ensuring your Jellyfin instance remains up to date and secure. ⏻

REHYDRATED

THE 14TH GEN i7 ITX BUILD

Just how much cooling do you need for the Core i7-14700K, asks *Zak Storey*

Here we have it folks, the last of Intel's 14th generation of processors: the Core i7-14700K. When we got hands-on with this sweet little number, we knew we had to do something a little more exciting than just your traditional review and build combo, so this time around we had one simple question to ask. With how aggressive Intel has been (and AMD, for that matter) with its voltage and clock-speed ramping over the last few generations, just how much cooling can you get away with for a top-tier processor on the latest platforms? Could you theoretically do away with a liquid cooler entirely, instead opting for something a little sleeker, a little smaller, and a little Noctua? That's exactly what we're here to find out.

The core of this build is centered around the Core i7-14700K, paired with Noctua's NH-L9x65 Chromax Black. The eagle-eyed among you will also spot a familiar case here in the form of Hydra's Mini ITX Chassis. This case is sadly no longer in production externally, but you can pick it up from a number of North American manufacturers (OCPC in particular) if you're on the hunt for something a little special. We first featured the Hydra ITX back in October 2020, paired with an even smaller cooler in the form of Noctua's NH-L9i Chromax Black, paired with the Core i5-10600K. So with a larger cooler and a beefier chip, can you still get decent performance out of a flagship product with an open-air chassis and Noctua's best and brightest low-profile kit? Let's find out, shall we?

INGREDIENTS

PART		PRICE
CPU	Intel Core i7-14700K	$402
Motherboard	MSI MPG Z790i Edge WiFi ITX	$310
CPU Cooler	Noctua NH-L9x65 Chromax Black	$70
RAM	32GB (2x16GB) XPG Lancer Blade RGB DDR5 @ 6,000	$110
SSD	1TB Western Digital Blue SN580 PCIe 4.0	$59
GPU	PNY Verto GeForce RTX 4070 12GB	$550
Case	Hydra Mini ITX Chassis Limited Edition	$100
PSU	Corsair SF750 80+ Platinum Gen 1	$170
TOTAL		**$1,771**

PRICES CORRECT AT THE TIME OF PUBLICATION

https://content.jwplatform.com/videos/MoKPVRiZ-u2lN49He.mp4
Please type this URL into your browser if the link is broken

XPG
XPG
[EDGE]
HYDRA

PICKING PARTS

CPU

INTEL CORE I7-14700K

Intel's Core i7-14700K is by far the more interesting of the CPUs launched with its 14th generation of products. Although it doesn't have the APO clout of the Core i9, it makes up for that in its expanded Efficiency cores. In fact, out of all three processors, it's the only one that's had any internal hardware changes made at all from a generational standpoint.

Of course, you still get those eight performance cores driving the bulk of in-game rendering and professional prowess that you'd expect, complete with hyperthreading, but you also get an additional four efficient cores, bumping the total count from eight to 12 here.

On top of that, it also gets the slight clock-speed bump we've seen with this generation, and gives you a total of 28 theoretical threads to play with. This is along with all of that 13th/14th gen connectivity support we've come to expect, including support for DDR5, and PCIe 5.0, of course. ***$402, www.intel.com***

Motherboard

MSI MPG Z790I EDGE WIFI ITX

As we're running an ITX chassis, we're going to need an ITX mobo, and there's no better board at our disposal right now than MSI's Z790i Edge. Its clean-cut aluminum design and crisp feature set make it a surefire pick for an ITX board in 2024. In fact, sadly there are only five Z790 ITX boards out there right now, and second to ASRock's Z790M-ITX WiFi, this is the cheapest board you can get (coming in tied with Gigabyte's Z790i Aorus Ultra). Still, it's a well-rounded board with a 10+1 power phase design, a solid M.2 heatsink, and PCIe 5.0 support for GPUs (sadly, there's no M.2 PCIe 5.0 support here, just PCIe 4.0), and it even supports DDR5 all the way up to 8,000 MT/s with the right kit. Rear I/O's no slouch either. Although it's somewhat lacking in USB ports given the size (with only four Type-A, and one Type-C total), you do get WiFi 6E, 2.5Gb Intel Ethernet, DisplayPort and HDMI out, and a Clear CMOS button as well, making it an excellent-value pick with a lot of modern support. *$310, www.msi.com*

SSD

1TB WESTERN DIGITAL SN580 M.2 PCIE 4.0 SSD

We're going for some cheap and cheerful kick-ass PCIe 4.0 storage in the form of Western Digital's 1TB SN580 M.2 SSD. This comes in at just $59 at the 1TB mark, but you can also pick it up anywhere between 250GB and 2TB.

At 1TB, you can expect sequential read and writes in the 4,100 MB/s window, and IOPS in the 600-750K capacity at 4KB. On top of that, WD has included a five-year limited warranty and 600TBW endurance limit, giving you some decent confidence.

It's not the quickest SSD, certainly not compared to some of the PCIe 5.0 SSDs (check out Centerfold on page 52 for such a drive), but for the money, it's a fantastic budget solution. If you don't fancy using it as an OS drive, it will easily do as a secondary drive for backups, games, and media. *$59, www.westerndigital.com*

GPU

PNY VERTO GEFORCE RTX 4070 12GB

We've dialed back the GPU for this build to something a little more respectable—and ideal for 1440p gaming—in the form of this excellent PNY RTX 4070 12GB GPU. At the time of writing it's available for $550, but with the advent of the Super cards this month, we expect the price to fall significantly in the coming months, making it a good budget pick if you're looking for a solid 1440p graphics card.

Unlike its more potent siblings, PNY's RTX 4070 doesn't rely on a 12VHPWR600 connector, only needing a single 8-pin PCIe power to get the most out of the card. Otherwise, you still get a lot of the RTX 4000 features we know and love, including some seriously aggressive clock speeds, 12GB of GDDR6X, and all that RTX and DLSS goodness that the cards are known for.

Don't expect this card to be a king at 4K, though; you can game at that resolution, but you're going to need to rely on DLSS to take the brunt of that workload off your shoulders. Without it, you're looking at the 30fps mark, which is fine for consoles, but less so if you're used to gaming on a PC at 60 fps. *$550, www.pny.com*

CPU Cooler

NOCTUA NH-L9X65 CHROMAX BLACK

Here's the other key piece of this puzzle: the Noctua NH-L9x65 cooler. This sleek little number is a low-profile 65mm tall, 95mm wide, quad-pipe air cooler, complete with Noctua's NF-A9x14 PWM fan. It's all black, with no RGB, no fuss, and has compatibility with pretty much every modern socket from the last eight years (LGA 115X and AM4 included).

It's that overall footprint that's the real king, however, with a total size of 95 x 95mm, and a weight of just 413g thanks to a combination of copper heat pipes and aluminum fins. It's one of the most potent low-profile coolers out there to date. As for the fan, it maxes out at 2,500rpm, with a max airflow of 57.5 m3/h and max acoustical noise of 23.6 dB(a).

But can it tackle the might of the Intel Core i7-14700K? That's the real question. *$70, www.noctua.at*

PSU

CORSAIR SF750 80+ PLATINUM GEN 1

An oldie but a goodie, Corsair's SF series of power supplies have long held a place in our hearts at *Maximum PC.* With some serious wattage behind the small form factors, and a slimmer design than other competing units (measuring at just 3.94 x 4.92 x 2.5 inches), this 750W is a perfect fit for our build, particularly as it will only draw 500W from the wall under full load.

That said, the SF750 Gen 1 is fairly old at this point, and one thing we wish we'd have pushed for is to pick up one of Corsair's newer units, such as the SF1000L (which features the newer, smaller, Type 5 cables albeit it's a slightly larger unit), or alternatively something like the Asus ROG Loki 1000W or Silverstone's SX1000-LPT. *$170, www.corsair.com*

CORSAIR
SF750
PERIPHERAL & SATA
24 PIN ATX
6+2 PCIe & 4+4 CPU

RAM

32GB (2X16GB) XPG LANCER BLADE RGB DDR5 @ 6000

XPG
XPG
DDR5

This issue, we've gone for a far more affordable set of RAM (you can read our full review on this kit on page 87), and that's ADATA's XPG line of Lancer Blade RGB DDR5. It's a low-profile kit that clocks in with a comfortable 32GB kit at 6,000 MT/s, with a CAS latency of 30, giving it a fairly comfortable 10ns real-world latency.

It's the price that really has our attention, however. Just $110 for a kit of DDR5 32GB is a staggeringly attractive offering, particularly when it has RGB, as well as performance metrics like that.

We've gone for the white model for this build, but it's also available in black. It's also worth noting that XPG has its own RGB software for lighting configuration (XPG Prime), which is fairly primitive in comparison to some of the other competition out there, but it will get the job done, and doesn't conflict with any other major manufacturer's work, in our experience.

$110, www.xpg.com

Case

HYDRA MINI ITX CHASSIS LIMITED EDITION

First up, we have some bad news. At the time of writing (and post-build shoot), Hydra as a brand is no longer an external consumer-facing company. That said, you can still pick up a number of Hydra cases from manufacturers across the US. OCPC sells it on Amazon, and Newegg's prices range from $80 to $100.

Stock is limited, so you may need to be quick if you want a piece of this chassis history for yourself. Hydra's ITX case is an open-air folded stainless steel chassis, complete with support for ITX boards, two-slot GPUs, and SFX PSUs. It's a beautiful little number that will be missed.

If you can't find a Hydra ITX for your build, there are still a number of ITX cases out there that will perform just as well, including Corsair's 2000D Airlfow, Hyte's Revolt 3, NZXT's H210i, and Fractal Design's Torrent Nano.

$100, www.newegg.com

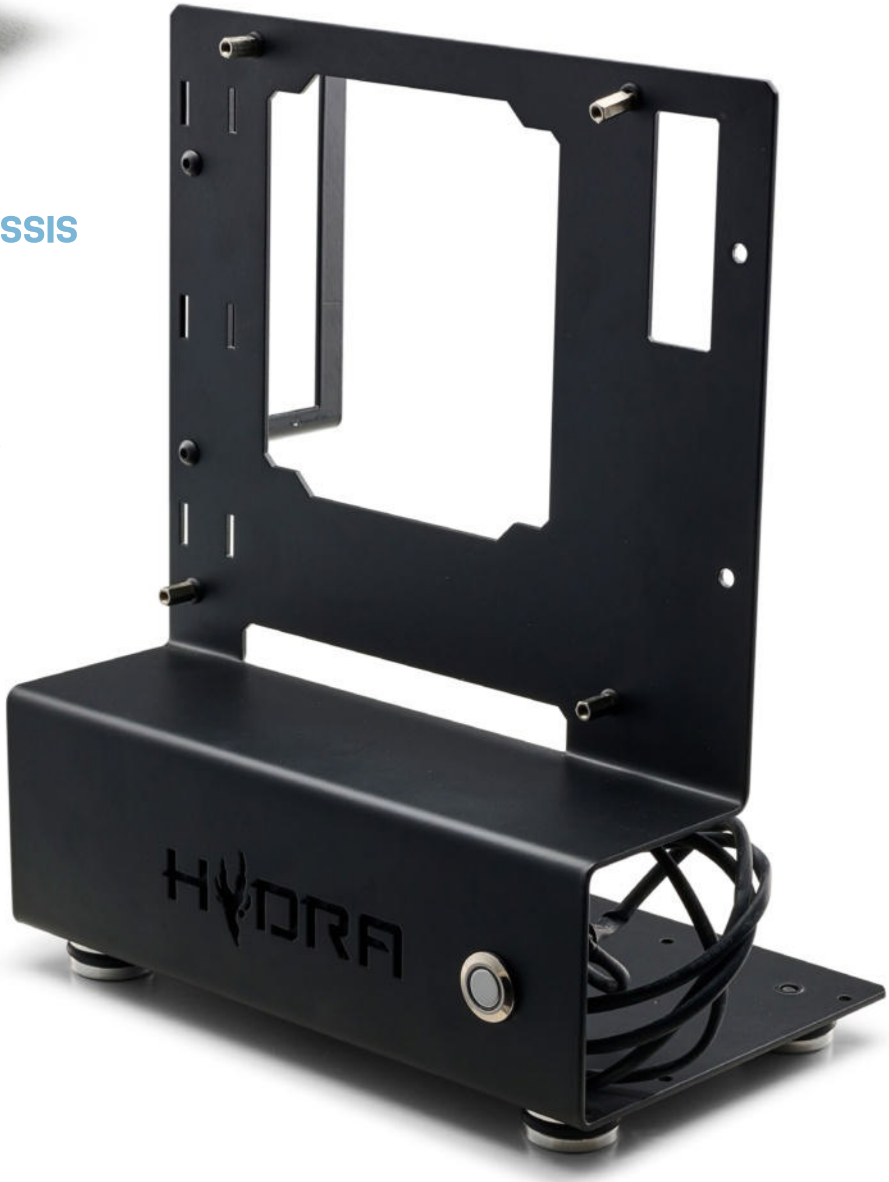

PUSH IT TO THE LIMIT

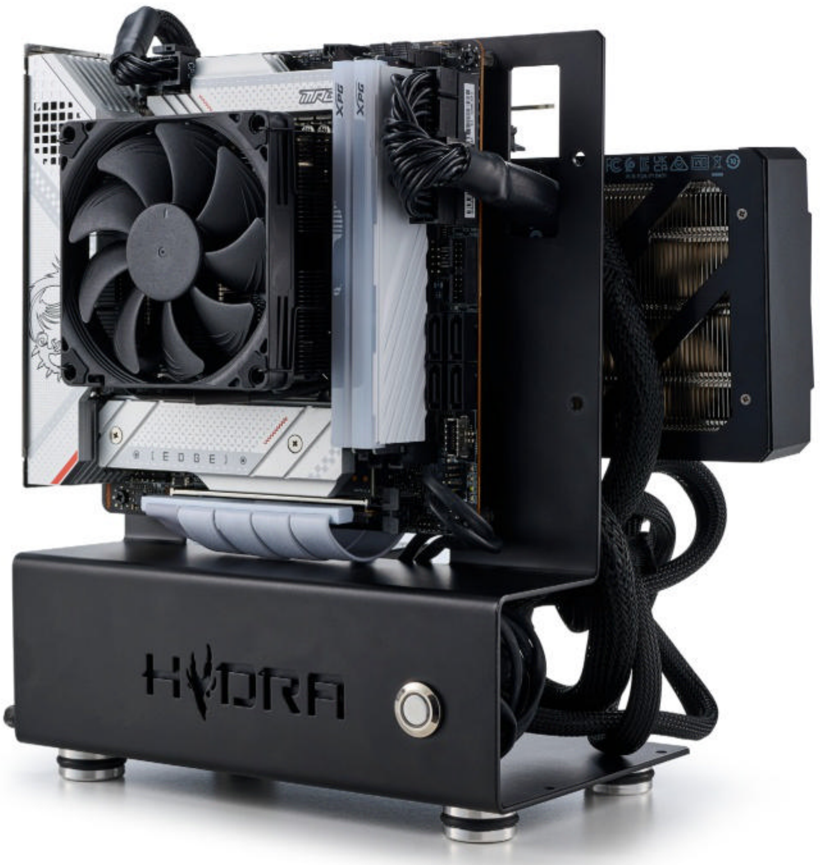

1

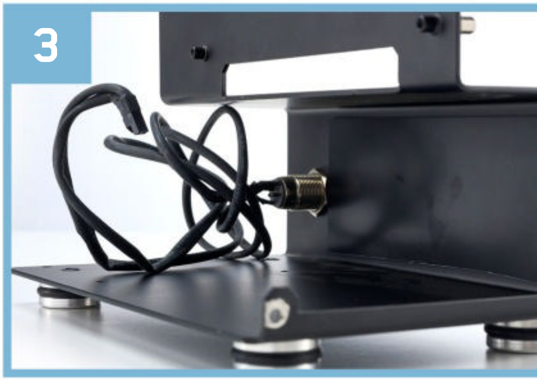
3

LENGTH OF TIME: 1-2 hours
DIFFICULTY: Medium

So then, with parts justified, let's talk shop. On the surface, this is a relatively simple build. Last time we used this Hydra chassis, we weren't entirely satisfied with the resulting build. There were a few foibles that arose after using the Hydra ITX, and during the build process that proved to be frustrating. Some of those issues stemmed from the design decisions made with the chassis, so those are two things we've adjusted and remedied from the get-go. On top of that, it's starting to age as a chassis a little, so we're going to need to do some minor cable upgrades just to keep up with the times, but more on that later.

The real focal point of this build is that Intel Core i7-14700K, and whether it can handle being cooled by a 65mm-high air cooler, because let's be clear, Noctua is very good at air cooling. They've long been our go-to for the best air coolers and fans around—second to none in the industry—but it's still a tall order to cool one of the hottest chips around with nothing more than a slim 95mm fan, four copper heat pipes, and a whole assortment of aluminum cooling fins on an open-air bench with nothing around it. Still, we're hopeful, and according to Noctua's own CPU compatibility list, it should be fully compatible, albeit with low turbo/overclocking headroom. Interestingly, Noctua actually has the 12600K on that list, with a high ranking for turbo/overclocking headroom, but jumping up the next-gen and its refresh sees that support plummet, which is a pretty good indicator of just how hard Intel is pushing this architecture.

That's something we've repeatedly seen with this series. It doesn't matter if you've got an Intel Core i5-14600K under a 420mm radiator, under load—it's still going to push its clock speeds and volts all the way up until it hits 100 C and its TJMax before throttling. On the one hand, that's darn impressive, and leads to some serious performance that's remarkably stable. On the other, do we really want processors and their architectures to be running so aggressively on a day-to-day basis? It's a tricky one to call.

A LIGHT AMOUNT OF MODDING

The Hydra ITX is an awesome case. There's not a whole lot out there quite like it. It combines a small form factor with an open-air design, and has a supremely simple build process due to how easy it is to access all the major areas and components. There's no excess cables, front I/O, or overly complicated cable management; just a single sheet of folded steel, mixed in with a number of mounting locations for hardware.

That said, it wasn't flawless, and in our initial build with it back in October 2020, two things stood out in particular. First up was the Hydra cutout in the PSU shroud area, namely the fact that you could see through it to the PSU itself. It's particularly annoying if your power supply has a white specs label on it. Secondly, the natural position of the GPU. In Hydra's installation manual, the way they have you configure that GPU shroud ensures you can keep your PCIe passthrough cable on the bottom of your graphics card, but because the case is reversed, it means that you end up with your GPU's I/O at the 'front' of the chassis, opposite the location

2

5

4

6

of your motherboard's Rear I/O. That's fine—you can get away with it—but it is mildly annoying if you want all your cables in one location on your desk.

We've tackled that final point first in a fairly simple manner. Simply put, the Hydra's small amount of self-assembly includes two Allen screws for mounting the GPU plate to the chassis itself. Interestingly, Hydra has also included mounting holes on the rear of the chassis. That means you can install the graphics card plate, reversed on the other side of the chassis, so the rear I/O for your GPU is flipped. This does mean your PCIe connector for your GPU will be at the top of the case (more on that later), and the PCIe power will be located on the bottom, but that's a sacrifice we're willing to make.

The first thing we've done is remove the plate [**Image 1**] and place it on the opposite side of the motherboard tray, effectively upside down [**Image 2**]. You'll notice that we've also added a couple of washers here to give us a bit of breathing room for inserting cables behind the graphics card itself. One thing we'll point out is that you don't need to worry about the fan/cooling situation or particularly chunky two-slot cards. Despite flipping the GPU bracket, your fans will still be facing outwards away from the case, so it will have plenty of breathing room.

PSU COVER BODGERY

To get around our transparent PSU cutout, we've gone for a very low-tech mod (because it's quick, easy, and looks the part). You can pick up some black plasticard from Amazon for about $5. This stuff is easy to cut and stick, and generally comes in a variety of colors and finishes. We've grabbed some of the matt black stuff they had in stock, roughly cut a piece to size with a pair of scissors, and then, using double-sided tape, stuck it to the inside of the PSU shroud [**Image 3**]. Flip over to the other side of the chassis and that PSU cutout, and you can see that the Hydra logo is now entirely blacked out [**Image 4**], but still looks just as classy as before. It's the small details that make a case personal to you, and this is one of them for us.

That said, it's entirely possible to go for something a little flashier here. Grab a piece of opaque white perspex, line it with an RGB LED strip, and you could have a glowing Hydra logo instead for a relatively low price. But heck, we like a clean, stealth build, and this one fits the bill just fine—all for relatively little outlay.

MOTHERBOARD PREP

With our chassis quibbles now satisfied, it's time to get onto the meat and potatoes of this build, and it's all going to begin with our motherboard prep. You've probably heard it a million times by this point, but it's always worth repeating. Wherever you're working, you need to do your best to avoid generating static electricity. Now, our hardware has come a long way in recent years. ESD protection is certainly something that manufacturers take into account, but the more you can minimize the risks, the better. In our case, we always recommend that you don't build a PC on a carpet, you don't wear wool, and always discharge yourself on a grounded source (plugged-in power supplies work well). Additionally, you're going to want to build and prep your hardware on a static-free surface.

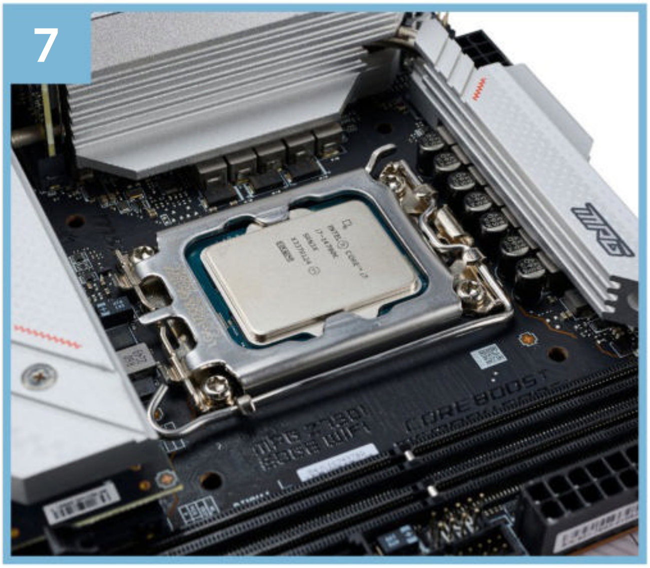
7

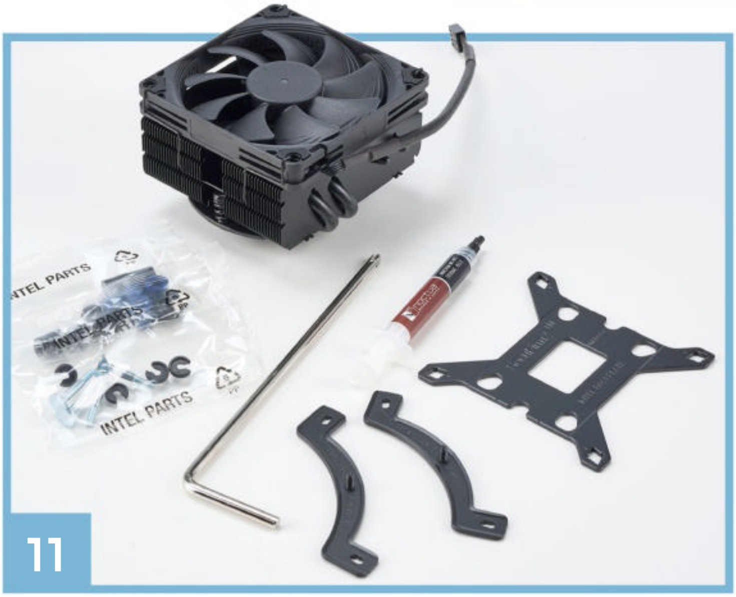
11

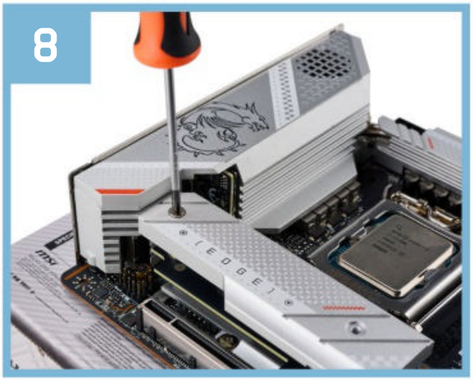
8

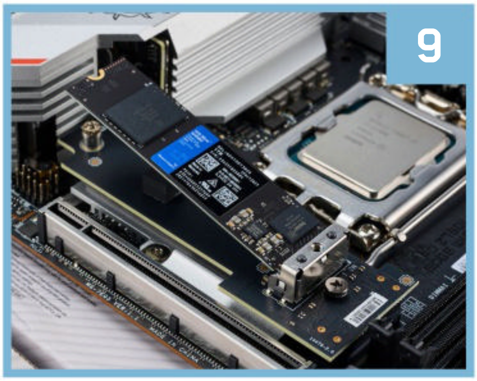
9

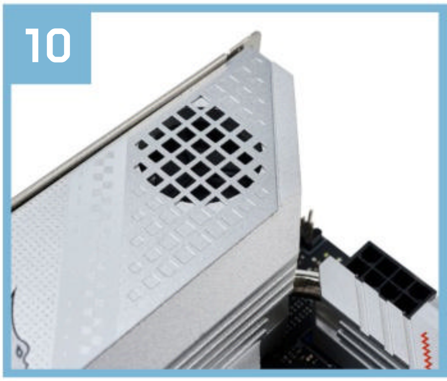
10

It's that last part that we're demonstrating here [**Image 5**], namely the outside of a motherboard box (or any product box, for that matter), which is a fantastic anti-static workbench. Interestingly, the anti-static bags that the components come in are only anti-static internally, not externally. The outside of the bag is lined with conductive material to help dissipate static away from the components internally, so the last thing you want to do is place your brand-new hardware on top of the bag. You might not kill it if you do, but the chance is never zero, and it's just not a risk worth taking.

INTEL'S LGA1700 SOCKET

Intel's latest socket is a bit of a beast. The LGA1700 completely re-oriented the design of the socket, and any who were familiar with any of the LGA1150 series or above might be shocked when they first see this thing. For a start, you've got a tighter bracket, a larger size, and the retention bracket itself lifts from the bottom rather than the top.

All other things are equal. To get access to the bracket, lift up the retention arm by moving it out and to the side. It will then ping up, giving you access to the primary retention bracket. Lift that up, and you'll be greeted by the socket underneath [**Image 6**]. At this point, you'll be ready to place your CPU into position, being careful not to knock or accidentally damage any of the pins.

Installing your LGA1700 CPU is a super-simple process. There are notches in both the socket itself and the CPU to ensure that you install it the right way. Line up the font so that the CPU name reads left to right, with the top left of the motherboard, and carefully line it up with the notches in the socket itself. Then, carefully place it in position, before doing the reverse of what you did earlier to unlock the retention bracket and arm. Once in and secure, you should be left with something that looks like this [**Image 7**].

M.2 TIME

More prep work. This time it's our 1TB WD SN580 SSD drive that we're going to be chucking into the motherboard. MSI's Z790i Edge comes with a heck of a heatsink and M.2 solution here—just below the CPU and between that singular PCIe slot.

You're going to need a small screwdriver to undo the screws holding the heatsink in position [**Image 8**]. Once removed, you can then install your M.2 SSD. Similar to our CPU earlier, your M.2 will also have a notch in it. Make sure you line that up with the notch in the M.2 slot and slide it into position. It will stick it out at an angle [**Image 9**] before you reattach the heatsink. That's completely normal, and ensures the M.2 makes good contact with the thermal pad once reinstalled.

Now it's time to re-attach the heatsink. Make sure you remove the protective film from the thermal pad underneath the heatsink, then carefully line the screw holes up with the mounts again and secure them back into position. After this, you'll be good to go.

Interesting side note: MSI has done some serious engineering work here to ensure its power solution and cooling is on point with this board. Take a quick glance at the top of the rear I/O [**Image 10**] cover on the Z790i Edge, and you'll spot this tiny fan that helps channel cool air

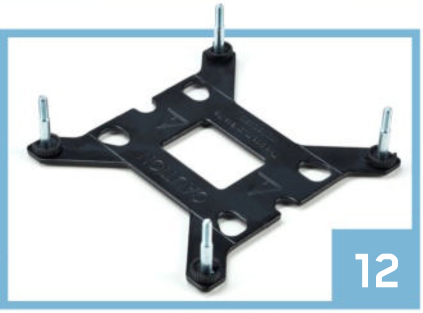
12

14

15

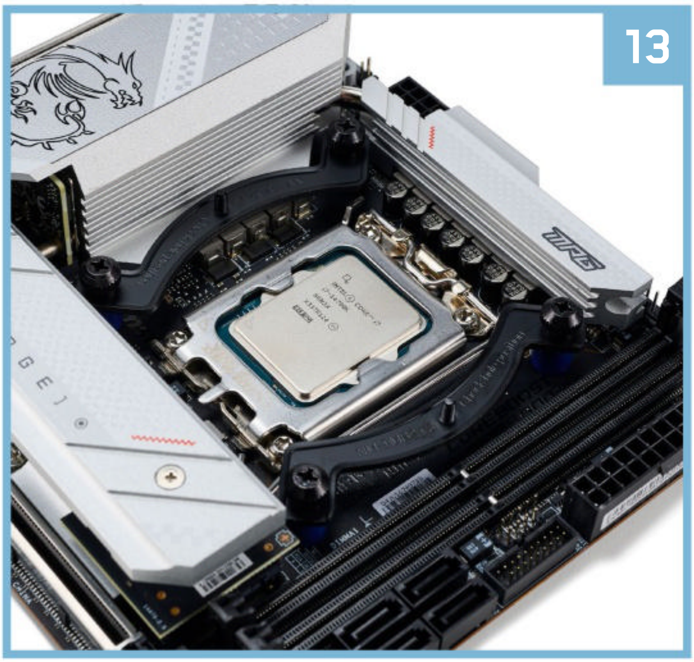
13

16

down and over the heat pipes, heatsinks, and power phases to keep everything running smoothly. Neat, huh?

BIG CHUNKY COOLING

Next up on the to-do list is installing the cooler. We're going to do this ahead of time so we don't have to worry about trying to fit this thing in around the RAM at a later date. Noctua's kits are always impressive feats of packaging, and the installation procedure is typically seamless. Take a quick look through the instruction manual and the parts included, and grab what you need. In our case, it's going to be all of the Intel parts, the screwdriver, the Noctua NT-H1 thermal paste (the best in the business, this stuff), plus the backplate and mounting kit, not to mention the cooler [**Image 11**].

Next, you're going to want to prep the backplate. Use the correct stand-off screws, and clip them into position with the included plastic clips [**Image 12**]. Once done, you'll want to carefully hold it in position through the motherboard socket, before placing the blue plastic spacers onto the standoff threads (there are different-color plastic spacers for different sockets—make sure you use the right ones, as LGA1700 does prefer a slightly lower CPU cooler for some cooler designs). Then, carefully place your mounting arms onto the threads, as seen here [**Image 13**], and temporarily secure them down with the thumb-screws by hand. Once secure, you can then tighten them off with a screwdriver. Do this by hand so that you don't overtighten them and accidentally thread the screws.

This is a snug fit, no doubt about it. You'll see at this point that there are two mounting screws facing upwards—these are the two points where your Noctua cooler is about to attach to. First, place a thin strip of thermal paste onto the middle of the CPU. You can chuck loads of this stuff on, or a small bead strip [**Image 14**]. It's better to have more than less, to ensure proper coverage. Noctua's NT-H1 is non-conductive, and although it will leave a mess if you ever replace the cooler, it's not going to impact or harm thermal conductivity in any major way (if you watch any competitive overclockers using LN2, you will see them pour entire tubes of the stuff on their CPUs before mounting their LN2 cylinders on top).

ON GOES THE BLOCK

Now, we're going to apply the block to it. You'll need to remove the fan first to access the screw locations. Line up the cooler with those two standoff screw threads we mentioned earlier, then carefully secure it down on both sides, a little bit at a time, until it's completely secure and the screw heads won't turn anymore [**Image 15**]. Once done, you're good to re-attach the fan.

Noctua's NF 95mm fan here is secured down via fan clips (if we're honest, our least favorite method). Simply connect them to each corner where appropriate, and secure them to the heatsink. Needle nose grips help here a lot. Before you do this, keep in mind the orientation of your PWM fan cable and the corresponding CPU fan header that you're going to need to attach it to [**Image 16**].

We've chosen to run our PWM fan header around the bottom and side of the cooler, wedged between the NH-L9x65 and the DDR5, and then up into the CPU

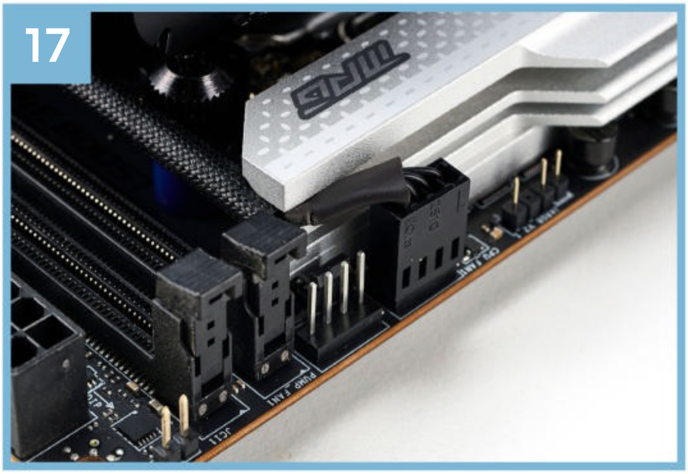

17

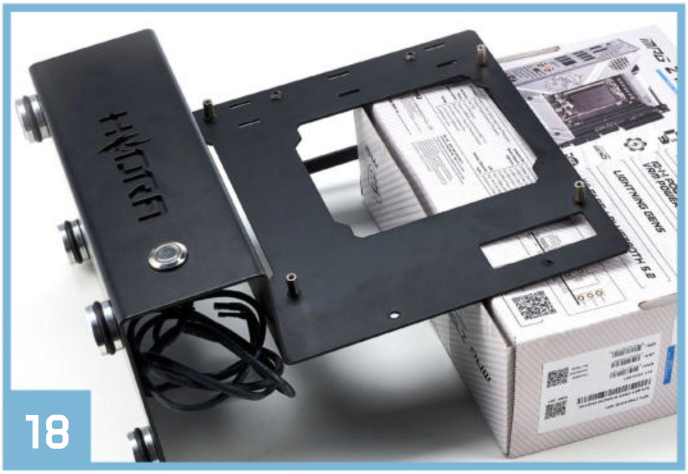

18

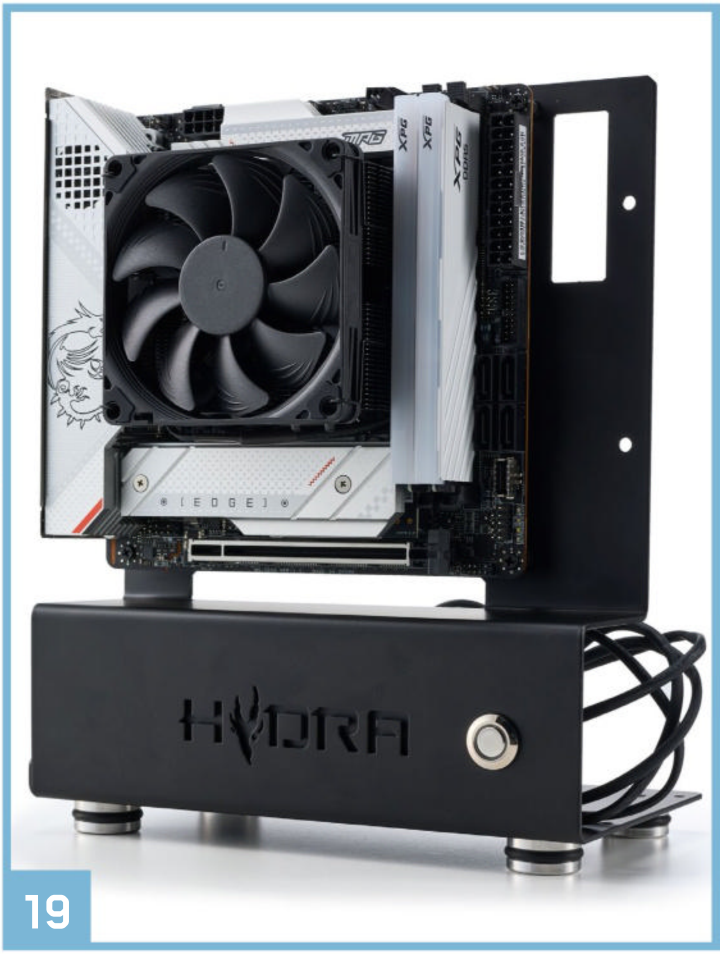

19

fan header, slightly wedged in the top-most heatsink [**Image 17**]. We've actually had to trim the shrinkwrap around the top of most of the cables to get this to fit properly. If you do need to do this, use a pair of scissors to carefully do so, making sure you don't accidentally cut any wires as you do it. It will give you a fractionally small amount of extra wiggle room to play with, but sometimes that's all it takes.

MOTHERBOARD INSTALL TIME

With the motherboard prep now mostly complete (we will pop the RAM in later), it's time to install it into the Hydra itself. Now, a unique 'feature' of the Hydra, is that thanks to its tiny size, open-air design, and GPU bracket orientation, it has a bad habit of falling over if you try to lay it on its side to mount hardware into it. It will just tilt sideways. To get around this, we recommend leaning it up against a box of some kind (we're using the Z790i Edge Motherboard box seen in [**Image 18**]). It's a tiny detail, but it will make your life so much easier.

With the chassis propped up, it's time to install the motherboard. Grab the motherboard standoff screws that are included with the case. Pop the motherboard into position, and then secure down each of the four corners. Once that's done, it's now a great time to install your DDR5 memory (remembering again to use the notches as an alignment guide), and you'll be set [**Image 19**].

Now is also a great time to install your front I/O for your power switch. Hydra uses a simple Dimastech Vandal Resistant Switch in its cases. They come with two pins attached: one for power, and one for power LED. Grab your motherboard manual and plug these into your motherboard's front I/O header. Here's where the fun begins. These cables aren't labeled, so there's no way of knowing which pin is for what. If you get to the end of your build, press the power and nothing happens, don't panic—check if your DimasTech switch pins are the wrong way around, before stripping the entire thing back to nothing (this is speaking from painful personal experience!).

INITIAL POWER SUPPLY INSTALL

With that done, it's now onto the power

PCIE RISER UPGRADES

As standard, the Hydra comes with a PCIe 3.0 riser cable for your build. It's a specific one built for the traditional setup that includes that originally oriented GPU facing the opposite direction to how we've built it here. Despite that, it still requires folding several times to fit in place and connect to the bottom of your graphics card.

In our case, given how old the PCIe 3.0 standard is now for GPUs, we've gone ahead and purchased a new unit from Amazon; LINKUP's Ultra PCIe 4.0 x16 Riser cable with a 90-degree socket in white. You can currently get them in a number of lengths depending on what you need, and they seem to be fairly solid (albeit expensive for what's effectively a cable).

Lian Li also sell a similar style cable for $90, but at 600mm long, which may

20

21

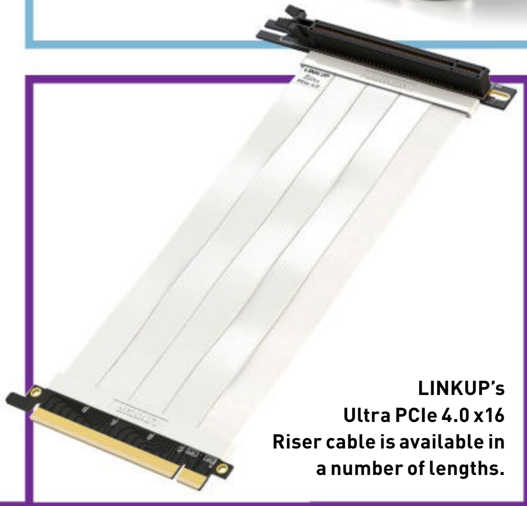
LINKUP's
Ultra PCIe 4.0 x16
Riser cable is available in
a number of lengths.

22

23

require some folding to make fit. One thing you'll also need to bear in mind is the orientation of the connections. PCIe cables (and graphics cards, for that matter) will only fit a certain way. We've specifically chosen the 90-degree socket. As we know, with the orientation of our GPU, this will plug directly into the card, without requiring us to fold it in any number of ways to get it to sit correctly, meaning we can keep it cleaner.

The big questions: do you need a PCIe 4.0 riser, and have we really saturated that bandwidth? It's tricky—for modern-day RTX 4000 series cards and top-tier 7000 series AMD GPU solutions, the answer is definitely. Anything pre-that era, that's a step up from PCIe 3.0 to 4.0, and it's a bit more difficult to discern the difference. In our testing on an older system we had in the Hydra, the 6900 XTX being used had horrendous coil whine with the PCIe 3.0 riser, so we still needed to buy a new riser.

© LINKUP

supply portion of the build. One of the advantages of this open-air chassis is that it's particularly easy to install your SFX PSU and the cables before or after the fact. Line up your fan so it's facing upwards (there's no cutout in the bottom, so you can't orient it in the other direction), then secure it in place with two of the included screws (we've gone with some spare thumbscrews that we had in the *Maximum PC* office, as it's a little easier to manage our older chassis with these as seen in [Image 20]). Next, you can install the power cables that you're going to need. In our case, that's the 24-pin ATX power, one 8-pin PCIe power, and one 8-pin EPS power for the CPU. We're also going to pre-route these as best we can, along with our PCIe 4.0 riser cable, to the top of the chassis in anticipation of our final step, the GPU install.

GRAPHICS CARD MOUNTING MISHAPS

At this point, you're going to want to install your graphics card in the machine. Turning the chassis on its side, with the motherboard facing down on a soft cloth, you should be able to slot your graphics card into the GPU bracket, attach the PCIe riser, and then secure it in place with the included Allen screw again.

What you don't want to do is try to secure it in place, making things mildly awkward because you've got an 8-pin EPS power cable stuck between the GPU and the slot, and then fully shear the screw head off (because this editor is clearly so incredibly strong). That's not what you want to do at all [Image 21].

Long term, the best solution would be to drill out the screw hole, re-tap it, then insert a new screw into its place, or ignore the tapping and use a bolt and nut instead. However, in our case, and because we were pushed for time, we had to cable-tie it into place [Image 22]. Let's be clear here—this is not a long-term solution; this is a temporary fix, and not something we would recommend anyone commit to [Image 23]. However, it does work for the time being, and certainly keeps the GPU securely in place.

With that done, and the PCIe power and riser cables connected, that's our build complete and ready to go.

MANY BUILDS ONE MONSTER

THE HYDRA IS, without a doubt, one heck of a unique chassis. Its folded steel design, unique aesthetic, and open-air style has ensured that it holds a special place in our hearts. The shearing off of that Allen screw in particular hurt, as it's a case that has lived in the *Maximum PC* offices since we first took delivery of it way back when. But it does prove that with some love and a few tweaks, an old chassis can still look the part and perform just as well as its latest counterparts.

There are a number of cases that fall into this category: Bitfenix's Prodigy line, NZXT's Manta, the Corsair Graphite 780t, the OG Fractal Design cases—you name it. They might not have the latest RGB or the best cable management, but if you can make them work, they will still serve you well, even in 2024.

All in all, the build process went surprisingly smoothly. Even with us snapping off that GPU screw, there were few hiccups or moments that challenged us. That's to be expected when your only cooling is a single CPU cooler.

Patching up those mild case foibles we weren't happy with has made a world of difference, but there's still room for improvement, particularly in the power supply department. Swapping out that older-style SF PSU for something a little sleeker, like the SFL series, and getting some of those braided Type 5 cables from another manufacturer would go a long way to making this case even better, shortening the cable mess in the back.

That is one thing we did forget to mention. To cable-manage this thing, you effectively have to bundle the cables together and tie them into loops as best you can behind the tray to hide them from sight. It is mildly difficult, and you'll still always be able to see them. Custom cables here would help, particularly if you could cut them to the right length. Additionally, the GPU still isn't supported as well as we'd like. Finding a better way to anchor that to the case at the other end would be a fantastic addition to this build, and make us feel a lot more comfortable in the long term.

That said, this is still an epically compact gaming PC build, perfect for LAN parties, traveling, or gaming on a big screen. In fact, fun side-note, instead of shipping this back for testing from the photo studio, we just chucked this inside a backpack, and carried it all the way home for testing without a problem.

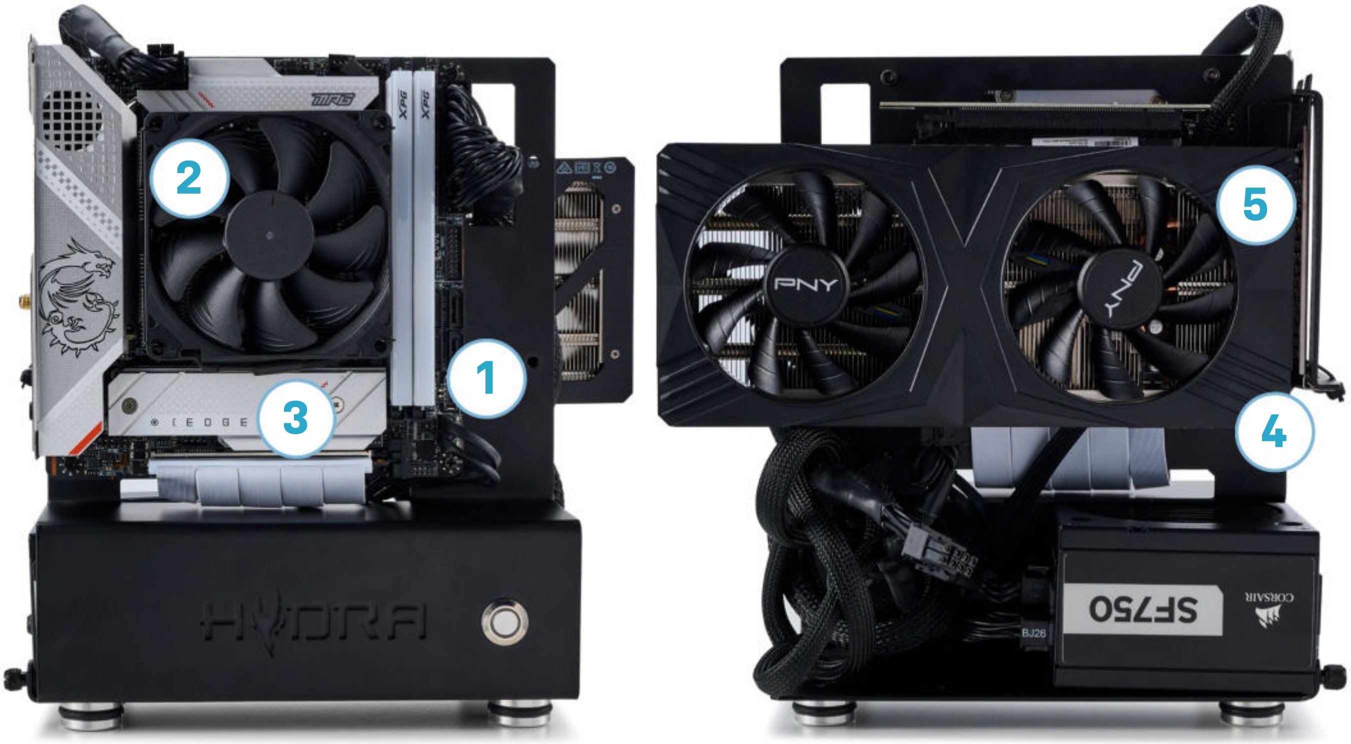

1 **It goes without saying that the stock PSU cables for this build, well, suck. That's the first thing we would change. Go for something like the Corsair SF1000L, along with a custom braided set of its type 5 cables. That would tidy this build up a treat.**

2 **The CPU cooler here is at the bleeding edge of what you can achieve with a low-profile heatsink, but you could choose to opt for a chunkier Noctua solution instead if you really wanted to live comfortably with that Core i7-14700K.**

3 **The only real downside to this particular motherboard is the fact that the M.2—which is located just below the CPU cooler—is only PCIe 4.0, rather than PCIe 5.0 compatible, meaning no modern super-fast SSDs for us.**

4 **Again, we need to reiterate that this cable tie here is only a temporary solution, not something that we would recommend having in place in the long term. Remove it, drill out the snapped screw, and replace it with a better solution.**

5 **Thanks to the Hydra's unique design, as long as the graphics card is a twin-slot design, you can fit any GPU you want in here—just make sure it fits on your desk. Replace the 4070 with a 4090 and a 1000W PSU, and you'd be on to a real winner.**

BENCHMARKS

	ZERO-POINT	
Cinebench R23 Single-Core (Index)	2,204	2,117 (-4%)
Cinebench R23 Multi-Core (Index)	36,815	29,446 (-20%)
CrystalDisk QD32 Sequential Read (MB/s)	7,138	4,176 (-41%)
CrystalDisk QD32 Sequential Write (MB/s)	6,299	4,135 (-34%)
3DMark Fire Strike Ultra (Index)	24,114	10,099 (-58%)
Cyberpunk 2077 (fps)	103	26 (-75%)
Cyberpunk 2077 RTX (fps)	69	17 (-75%)
Metro Exodus (fps)	133	43 (-68%)
Metro Exodus RTX (fps)	109	30 (-72%)
Total War: Three Kingdoms (fps)	95	42 (-56%)
Core Price	$3,068	$1,431 (-53%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of the 14th gen Gaming PC from our last issue, featuring an Intel Core i9-14900K, Nvidia GeForce RTX 4090, Asus ROG Maximus Z790 Dark Hero Motherboard, 32GB (2x16GB) of Corsair Vengeance DDR5-4800, and 1TB Adata Legend 960 Max PCIe 4.0 M.2 SSD. All games tested at 4K "Ultra" graphics presets with DLSS and V-sync turned off and XMP for RAM speed turned on. No manual CPU overclocking. "Core Price" refers to the key components generating performance (CPU, GPU, Mobo, SSD, RAM), not accessories.

PERFORMANCE DIFFERENCES

LET'S BE CLEAR, our zero-point graphs above provides a bit of an unfair comparison. The RTX 4090 featured in it is an absolute monster, and clearly demolishes the RTX 4070 when it comes to 4K gaming performance. Those figures aren't wrong here. You're talking demonstrable drops of up to 75 percent in pretty much all titles, particularly those with RTX features. *Total War: Three Kingdoms* was the closest fight between the two, clocking in at 42 fps versus 95.

It's also worth bearing in mind that the overall hardware cost difference between these two was over double, so that is to be expected to a lesser extent. Overall performance, though, with that cooler and the 14700K, was pretty solid. Single-core performance took a bit of a hit, unsurprisingly as the turbo clocks couldn't be maintained for as long as we'd have liked, but still, a drop of four percent isn't exactly mind-blowing. The 20 percent drop in multi-core performance, however, is a little more jarring, although the 14900K does have 25 percent more efficiency cores than the 14700K, and both chips feature eight performance cores, albeit the 14700K clocks in at 5.6 GHz on those versus the 14900K's 6 GHz at peak.

What we found in our temperature testing was more surprising. In fact, pretty much every chip that we've ran through consecutive Cinebench R23 runs has topped out at 100 C on our test beds, its TJ Max. It doesn't matter if that's the Core i5-14600K under a 420mm AIO, or a Core i7 14700K under this little number, Intel seems to be pushing these processors to the limit, keeping them stable, and drawing out as much clock speed as it can, temperature be damned.

That's an interesting position for the processors to be in, because as we can see here, despite Intel recommending the 14700K and 14900K sit under an AIO, they can operate fairly well with a 'lower' spec air cooler.

Ultimately then, this rig is a solid countenance, and you can in fact, comfortably run the Intel Core i7-14700K under a 65mm tall air cooler without major hindrance. If all you do is game, and you're not interested in the world of cutting-edge video rendering, going the budget route and opting for a far more affordable system might be a better bet. Save the cash on the cooling and E-ATX motherboards, and instead reinvesting that into your GPU might be the best bet.

DISCOVER THE ALTERNATIVE TO APPLE AND WINDOWS

This is a treasure-trove of Linux and open source knowledge, from the evolution of Ubuntu and its mobile platform, to projects like automating your home with Raspberry Pi.

NEW
LINUX
2024 ANNUAL FORMAT
» HACK YOUR GRAPHICS
» GET INTO CODING
» ESCAPE WINDOWS
» FREE ALL YOUR FILES
AND MUCH MORE!
OVER 160 PAGES OF EXPERT TUTORIALS & MUST-READ FEATURES
ALL THE BEST CONTENT FROM THE NO.1 LINUX MAGAZINE
ISBN 978-1-80521-141-9
Linux Format Annual 2024
VOLUME 7
ON SALE NOW

FUTURE

Ordering is easy. Go online at:

magazinesdirect.com

Or get it from selected supermarkets & newsagents

INTEL'S 14TH GEN, A RECAP

RAPTOR LAKE IS OVER, WHAT'S NEXT?

INTEL'S RAPTOR LAKE chips have been an interesting generation compared to some of the iterative updates we've seen from Team Blue over the years. The 14th gen in particular hasn't exactly knocked the socks off the team here at *Maximum PC*.

Comparatively, the 14600K only had a clock speed bump compared to last gen, the 14900K got a clock-speed bump and access to APO (application optimization), which was fairly underwhelming at launch, and the 14700K, the only chip to receive a clock speed increase plus significant hardware changes (namely four additional efficient cores), has only helped to blur the lines between the Core i7 and the Core i9 even further, particularly given APO's lackluster arrival.

So what is the strategy from Team Blue? Well, it's hard to tell. This generation does feel like a marketing swing, more so than a generational launch. Performance has been top-tier, as we'd expect from processors at the cutting edge of Intel's consumer product stack. In fact, in pretty much all benchmarks we can throw at them, all of these chips have produced some fine results, but it doesn't feel like a generational shift or a typical launch. There have no been no i3s, no non-k variants, or even special editions. Just three processors, all of which are fairly forgettable in contrast to Intel's 13th gen.

Noctua's own CPU compatibility chart shows the differences here best. Jumping from the 12600K to the 13600K saw far greater demands on the cooling front to the point that some of its own flagship products could no longer comfortably keep the chips performing at their top turbos for any long period of time. In fact, in our anecdotal testing, and as we've already mentioned in this very build, products like the Core i5 to the Core i9 frequently max out at 100 C under load like it's nothing. Efficiency seems to have been thrown to the wayside to leverage a win, and to be frank, it's working. Raptor Lake's process is clearly now so mature that rather than leave these processors out there to the silicon lottery, Intel's aim has clearly been to shift these new, superior chips, into a generation of their own.

It's logical, generates marketing hype, cements its processors as top-tier products still, cleans out the old stock, and allows them to keep that regular product launch schedule in check, as it prepares for the move to LGA 1851 and Arrow Lake/Lunar Lake in 2024.

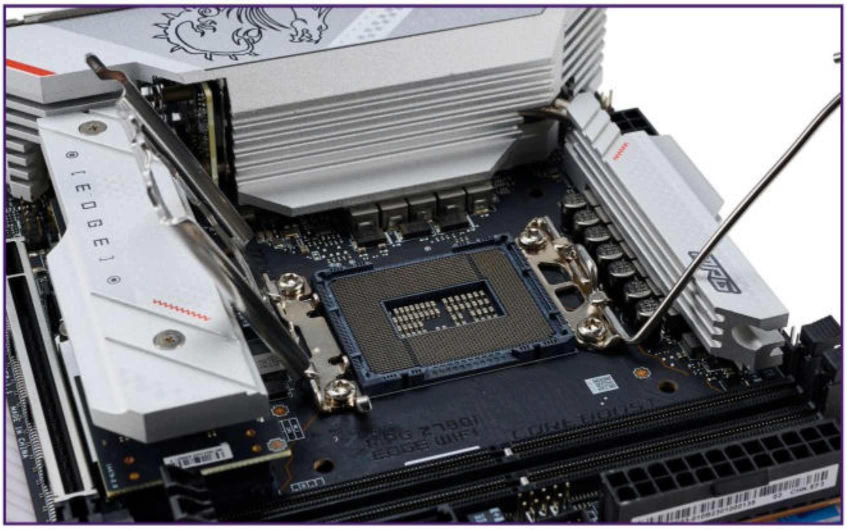

Intel's next-generation processors will fit in the LGA1851 socket, and the processors and socket will be the same size, so current coolers will fit, but there'll be 151 extra contacts.

LGA 1851 AND ARROW LAKE

Understandably, details on the 15th generation of Intel processors are light on the ground. We don't know for sure when they're landing—all we know is that the socket is going to be called LGA1851 (with 1,851 contacts instead of the 1700 we currently have). There's a rumor that there's going to be additional PCIe 4.0 lanes, and we'll lose access to DDR4 support on these new generations, and of course that the 'Core' branding moniker is going the way of the dodo, but that's it.

We expect Intel's going to bank heavily on APO moving forward and lean more into AI and machine learning for those kinds of applications, but outside of that, there's little to go on. Its performance-core and efficient-core chip design is holding up well, and logically makes sense. There's a chance that they could go the 3D-stacking route in a similar manner to how AMD has with its 3D VCache, but the big concern there will be heat, and how to manage that, particularly given how power-hungry these already are.

Intel's 14th gen has been a bit of a mixed bag as far as launches go.

INTEL CORE™ i9
i9-14900K
SRN48
X330K254

SHOULD I BUY A 14TH GEN PROCESSOR?

This is the big question that needs answering, and the answer is that it's going to depend on your circumstances. If you're already on 13th gen, or even 12th, for that matter, the likely answer is 'no'. While performance is impeccable, from a generational standpoint there's little advantage to spending the additional cash.

The only exception to this rule is if you're looking to upgrade your graphics card as well, or you need faster storage to take advantage of that PCIe 5.0 specification. If that's the case, and you need to be at the absolute cutting edge, then potentially it does make a lot more sense to jump ship.

That said, we expect that Intel may very well launch a new series of processors either at Computex in June, or potentially even earlier in the year, in time for the September quarter. In this case, it may very well be better to hold off until then.

Boosting frame rates via AI and clever programming

NVIDIA HAS BEEN pushing so-called 'neural rendering' techniques since the launch of DLSS in 2018. While DLSS had a bit of a slow burn at launch, there are now more than 500 games and applications that use Nvidia RTX features.

The core idea of neural rendering is to leverage AI models to improve the quality and performance of games and other graphics applications. As pixels become increasingly complex to render, figuring out ways to reduce the number of fully rendered pixels and then interpolating to fill in the gaps can provide a better overall experience.

However, Nvidia's solutions are designed to only work on Nvidia GPUs. Enter teams red and blue with alternatives that can work on a wider set of hardware. Upscaling and frame generation are here to stay, but how do the various AMD, Intel, and Nvidia solutions stack up, and what does the future hold for neural rendering techniques?

Join us as we cover the state of the upscaling industry and related technologies. **–JARRED WALTON**

UPSCALING SHOWDOWN

DLSS, FSR, & XESS

© EPIC GAMES

AHMA
AHMA
LIMU
Lime-Lemon

Alan Wake 2

Nvidia uses the umbrella term 'neural rendering' for its AI-based DLSS features that aim to boost frame rates by reducing the number of fully rendered pixels your GPU has to generate.

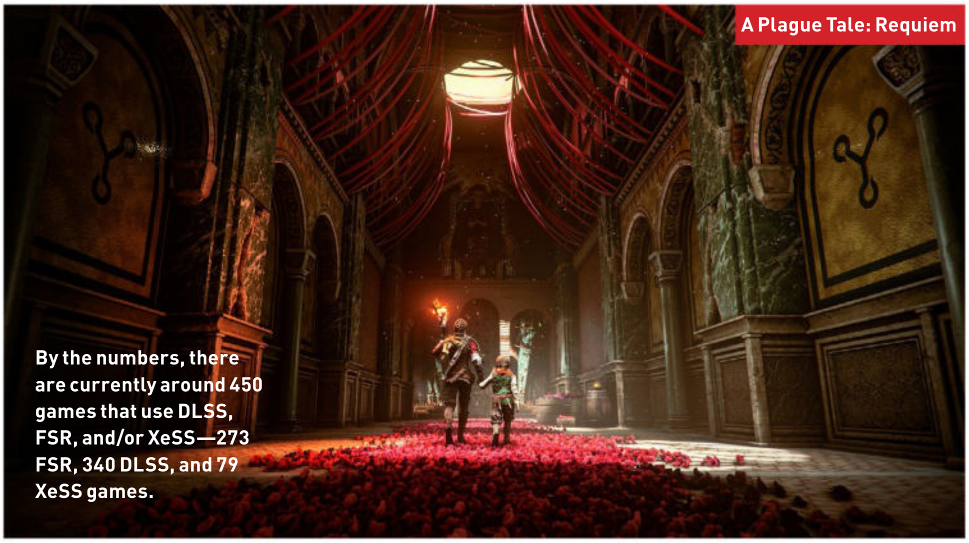

By the numbers, there are currently around 450 games that use DLSS, FSR, and/or XeSS—273 FSR, 340 DLSS, and 79 XeSS games.

UPSCALING 101: THE ALGORITHMS

Fundamentally, upscaling isn't a new idea. From the very first 2D sprite, games have been using upscaling algorithms. More recently, real-time upsampling of video content became an important feature, and we've seen various solutions on DVD, Blu-ray, and HDTV devices over the past couple of decades. Even before DLSS arrived, upscaling was available in games. All you need to do is run a game at a lower resolution than your display's native resolution, and some form of upscaling happens, either via the GPU or the monitor.

But we're more interested in the modern upscaling algorithms in games. At present, the three contenders are Nvidia DLSS, AMD FSR, and Intel XeSS, but there are different versions of each of those, with later iterations generally providing improved quality and additional features like frame generation. Let's quickly cover the basics of the three solutions, starting with Nvidia DLSS.

NVIDIA DEEP LEARNING SUPER SAMPLING

Nvidia DLSS launched in 2018 as a spatial upscaling algorithm. It required game-specific training on Nvidia's supercomputers, with tens of thousands of images, provided in pairs of lower-resolution 'input' and high-resolution 'ground truth' quality outputs. These would feed a deep learning algorithm that would be trained to create higher quality outputs given lower-resolution inputs. A key part of the algorithm is that it handles both upscaling and anti-aliasing—the removal of 'jaggies' on high-contrast edges.

It all sounded nice in theory, but it proved to be less than ideal in practice. Only a handful of games ever bothered to try to implement DLSS 1.x, and several of those would later get DLSS 2.x upgrades. There were multiple issues. First, all the per-game training added complexity—game developers couldn't simply plug in a working solution; they had to capture pairs of frames and send those to Nvidia's supercomputer. Second, the quality was lacking, with perceivable blurriness. Finally, the algorithm wasn't very flexible, so games like *Battlefield V*, for example, would only allow the use of DLSS at 4K on an RTX 2080 Ti—if you were playing at 1440p, the option was locked out. That was also because the spatial algorithm didn't scale to higher FPS very well, so upscaling to a target 1080p could result in a reduction in fps.

In a massive overhaul that came out early 2020, DLSS 2.0 ditched the spatial upscaling algorithm and switched to temporal upscaling. That provided a lot more data for DLSS to work with, as it now receives the current and previous frames, plus motion vectors and depth buffers. It also became a generalized algorithm, so it was no longer necessary to train DLSS 2.0 for each game on new images, and the result was far more flexibility.

This marked the point where DLSS went from being 'nice in theory' to being something that was useful in practice for most situations. There were still bugs to iron out, and over the past three years, DLSS 2.x upscaling has continued to improve in quality. Improvements focused on the elimination of ghosting (repeating parts of an image that shouldn't occur), as well as improvements to overall image quality, thanks to additional training.

DLSS 3 was introduced with the RTX 40-series hardware in late 2022. While the name might suggest further enhancements to the same base algorithm, DLSS 3 doesn't have anything to do with 'super sampling', and is instead focused on frame generation. Dedicated hardware on the 40-series GPUs, the Optical Flow Accelerator (OFA), takes two different frames and interpolates an intermediate frame, potentially doubling the perceived frame rate of a game. DLSS 2 upscaling is required as part of the package, as is Nvidia Reflex—a solution to reduce

RTX 4070
RADEON
ARC A770
Limited Edition

While you need an Nvidia GPU to utilize the company's DLSS tech, AMD's FSR and Intel's XeSS are open to be used by rivals.

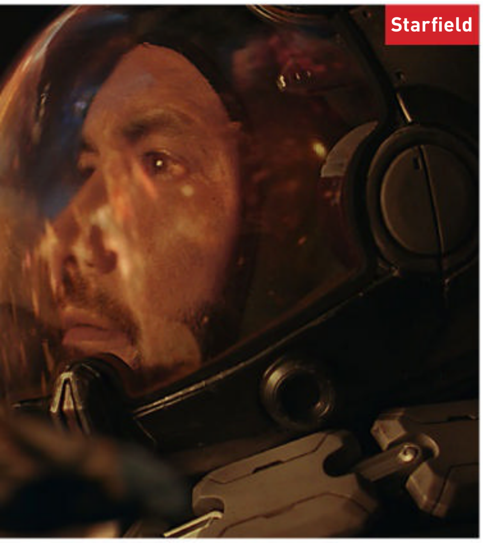
Starfield

The Ascent

input latency. Only the latest RTX 40-series GPUs support DLSS 3, though earlier cards can still make use of upscaling and Reflex.

This is where things get confusing, because now there's DLSS 3.5, aka ray reconstruction. Unlike frame generation, ray reconstruction works on all RTX GPUs. However, this is an AI-based denoising algorithm focused on improving ray tracing quality, which makes it less practical to use on old RTX 20-series cards, though it's still possible to try it. Only a couple of games currently support DLSS 3.5, however—*Cyberpunk 2077* and *Alan Wake 2*—and it's only available when full ray tracing is used.

AMD FIDELITYFX SUPER RESOLUTION

AMD countered the DLSS onslaught with FSR in 2021, a 'universal' upscaling solution that doesn't require any special AI hardware or training of an algorithm on supercomputers. FSR works on nearly all GPUs, including those from Nvidia and Intel, and AMD has been shadowing Nvidia with new variants and enhancements over the years, challenging Nvidia's proprietary AI solutions.

The original FSR used a spatial upscaling solution. It was easy to implement, and thus found its way into quite a few games. However, the overall image fidelity was lacking. FSR 1.0 was basically Lanczos upscaling with additional edge detection and anti-aliasing code integrated to try to improve the output quality. It's worth pointing out that Nvidia has offered an alternative Lanczos sharpening filter in its drivers since 2016.

FSR 2.0 was a complete reworking of the core algorithm, and this time it required much tighter integration into a game. It's now a temporal upscaling algorithm, and like DLSS 2, FSR 2 requires multiple inputs: the current (lower resolution) frame, the previous frame, the depth buffer, and a motion vector buffer. Quality improved quite a bit, though FSR 2 requires significantly more processing time than FSR 1.

Most recently, AMD released FSR 3, which now has frame generation support. It also requires the integration of support for AMD's Anti-Lag+ technology to reduce the added latency created by frame generation. Of course, Anti-Lag+ only works with AMD

VARYING FSR 3 IMPRESSIONS

Frontiers of Pandora currently represents one of the best implementations of AMD's FSR 3 with frame generation we've seen.

FSR 3 right now is very much in its infancy. At present, there are only four games shipping with FSR 3 frame generation. We've tried three of them, and had very different experiences with the end results.

***Forspoken* and *Immortals of Aveum* were first out the door, and felt a bit raw—turning on FSR 3 frame generation would boost average fps, but often the minimum fps wouldn't change, and the games would still feel choppy. There were also issues with having vsync disabled, or with using variable refresh rates alongside frame generation. Initial impressions were decidedly lackluster.**

Then *Avatar: Frontiers of Pandora* came out, and FSR 3 frame generation seemed to work much better. In fact, unlike most games, even generated results of 40 fps remained playable. This flies in the face of how frame generation is supposed to work, but it applied to a variety of GPUs. Running around Pandora on an RTX 4060 and averaging 45 fps at 4K ultra (with FSR 3 quality upscaling and frame generation) was totally viable.

Is there something else going on under the hood? We don't know, but FSR 3 source code is available, and perhaps the developers were able to tweak the algorithm to improve the overall feel. With such a small sample size, the jury is still out, but FSR 3 works on all GPUs, including Nvidia's RTX 30-series parts that don't have access to DLSS 3. In our book, that's a big potential draw.

© NVIDIA, FOCUS ENTERTAINMENT, BETHESDA SOFTWORKS, CURVE DIGITAL, UBISOFT

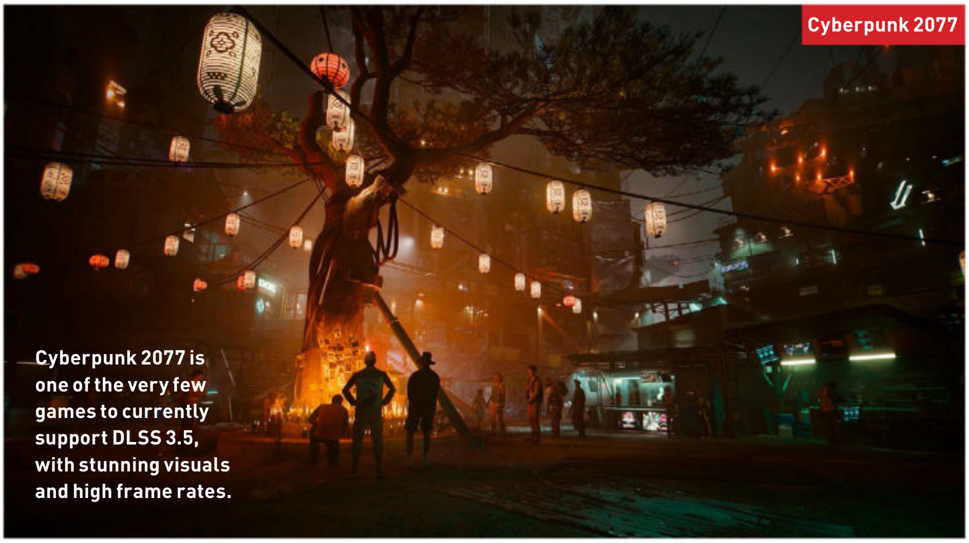

Cyberpunk 2077 is one of the very few games to currently support DLSS 3.5, with stunning visuals and high frame rates.

GPUs. As with DLSS 3, games that list FSR 3 support include upscaling as well frame generation.

At present, there are over 200 games with FSR integration, but more importantly, there are over 140 games with FSR 2.0 or later support, with another 50 or so upcoming games that will support either FSR 2 or FSR 3. Native support for FSR 3 remains quite limited, however, with only four currently shipping games offering frame generation: *Avatar: Frontiers of Pandora*, *Forspoken*, *Immortals of Aveum*, and *Like a Dragon Gaiden*.

INTEL XE SUPER RESOLUTION

Not to be left out, Intel's XeSS arrived in late 2022 alongside its dedicated Arc graphics cards. It sort of straddles the lines between DLSS and FSR, with several modes of operation depending on your GPU. The base algorithm uses a deep learning network, similar in many ways to Nvidia's DLSS. The catch is that you can use XeSS on non-Intel GPUs, though it may not have quite the same level of quality or performance in that case.

XeSS uses a temporal upscaling algorithm—Intel skipped the spatial upscaling, likely as it was proven mostly unnecessary. It takes the same three inputs as the other solution: frame buffer, previous frame, depth buffer, and motion vectors. The trained algorithm then uses those to upscale to a higher-quality output.

XeSS operates in one of three modes. The highest performance and quality mode requires Intel XMX units, which are found on dedicated Arc GPUs, but not on the new integrated Arc GPUs in Intel's latest Meteor Lake processors. These are matrix units similar to Nvidia's tensor cores, designed to accelerate AI workloads. If a GPU doesn't have XMX units, XeSS will run in DP4a mode. DP4a instructions are 8-bit integer shader instructions that are supported by modern GPUs, designed to help accelerate AI workloads. Therein lies the catch.

First, DP4a won't have as much raw computational power as XMX, so the underlying algorithm tends to produce lower-quality outputs. However, more recent versions of XeSS (1.1 and later) have improved the overall quality of the DP4a path. Also, there

MIX AND MATCH?

Many games support multiple upscaling algorithms, like *Modern Warfare 3*, but no more than one at a time.

Here's a thought: what if you have an Nvidia card and you enable DLSS upscaling and frame generation alongside FSR 3 upscaling and frame generation? There have been a few hacks that attempted to do just that, but in general this isn't allowed in the game engines.

Upscaling an already upscaled frame would be like making a photocopy of a photocopy. Each extra level of copying would become worse. But it's more than just the potential for artifacts. Game engines pass the frame buffer, previous frame, depth buffer, and motion vectors to the upscaling algorithms.

Most of those buffers would be at the lower rendered resolution, but the previous frame would be at the final output resolution. If you wanted to do upscaling twice, the engine would need to also spatially upscale the depth buffer and motion vectors to pass to the second upscaling algorithm.

Upscaling of upscaling would compound the artifacts, and isn't likely to look good. But it's theoretically possible, and it would be fun for games to allow users to try it. But for now, the games we tried all force you to opt for one of the available algorithms.

Many games also lock you into upscaling and frame generation from the same core algorithm, so you can't use DLSS upscaling with FSR frame generation. There may be exceptions, but running both APIs at the same time is certain to introduce additional overhead.

© CD PROJEKT RED, ??????, ACTIVISION

Immortals of Aveum

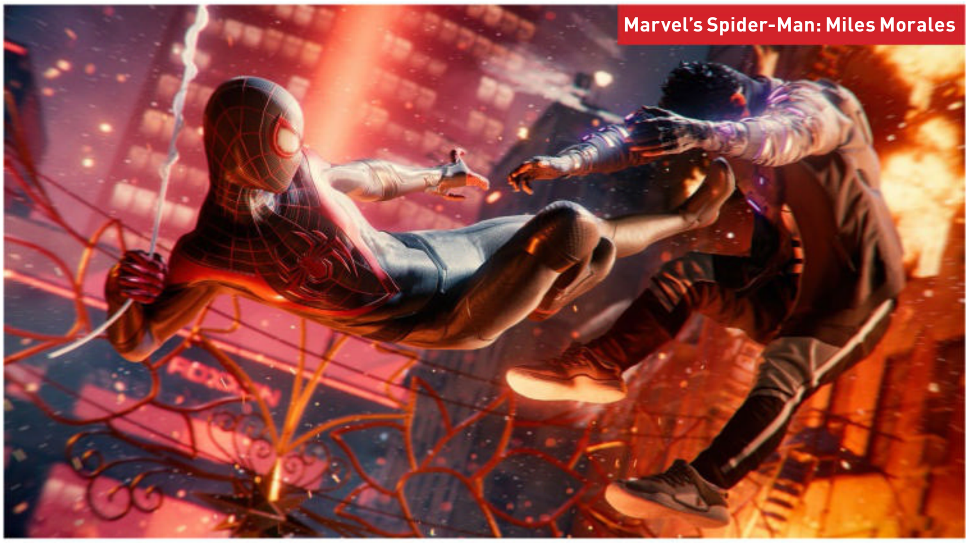
Marvel's Spider-Man: Miles Morales

are two DP4a code paths, one optimized for Intel GPUs—from the Gen10 graphics in Ice Lake processors up to modern Meteor Lake Arc iGPUs—and the other designed for non-Intel GPUs.

XeSS represents the Johnny come lately upscaling solution, and as such it's in the fewest number of games: 77 total at the time of writing, with 27 running XeSS 1.1 or later. There are only a few games that only support XeSS, with most offering either FSR or DLSS as an alternative. That makes sense, as once a game engine has added support for any of the three temporal upscaling algorithms, it should be relatively easy to integrate support for the others as well. Considering that Intel Arc GPUs currently only account for a small percentage of the total GPU market, most games tend to focus first on FSR or DLSS integration.

UPSCALING QUALITY COMPARISONS

While there are hundreds of games that support upscaling using one of the three major algorithms, there aren't that many games that opt to support all three solutions—just over 40, by our count. It's messier than that implies as well, because a game that supports DLSS, FSR, and XeSS doesn't necessarily mean all three received equivalent effort, or that all three are using the latest and highest-quality versions of the upscalers.

The quality of upscaling and the performance advantages it might bring can also vary substantially across game engines. Some games are inherently more CPU limited, but we've also encountered games where there appear to be other bottlenecks. Image quality in general tends to be good for the most recent versions of each upscaler—though not necessarily perfect.

If you're capturing screenshots and videos to compare quality, there are relatively few cases where we'd suggest that upscaling looks as good as or better than native rendering. Usually, the bigger issue is that native rendering may use overly strong TAA that blurs everything out, and switching to DLSS, FSR, or XeSS would replace the TAA with a less blurry algorithm. The upscalers can also include a sharpening filter to undo some of the blurriness, which often improves the overall image quality.

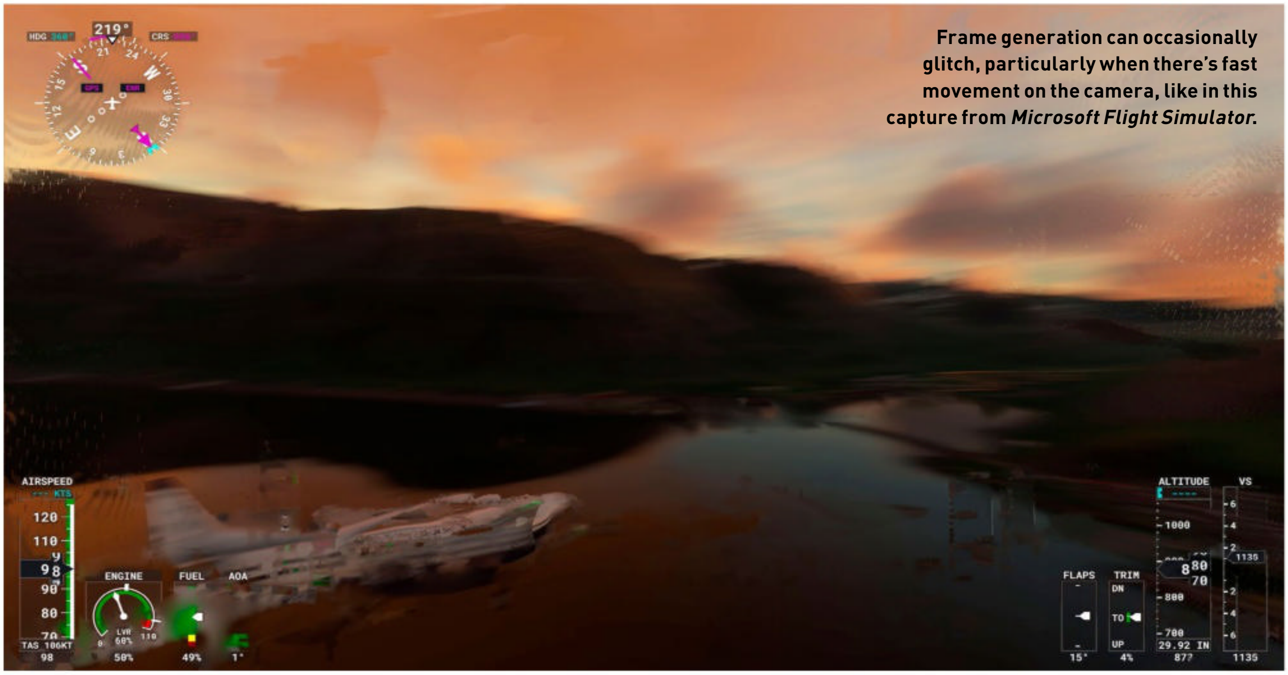

Frame generation can occasionally glitch, particularly when there's fast movement on the camera, like in this capture from *Microsoft Flight Simulator*.

© NICROSOFT, MARVEL

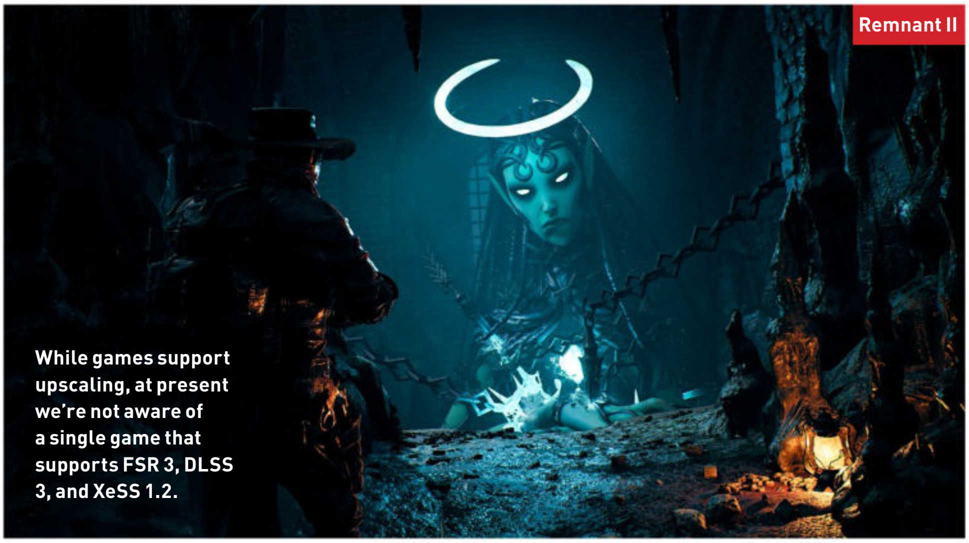

While games support upscaling, at present we're not aware of a single game that supports FSR 3, DLSS 3, and XeSS 1.2.

In motion, playing the games rather than trying to spot minor differences in the output means that things become a lot more nebulous. Occasional artifacts caused by upscaling become easier to miss when you're just having fun. That's especially true when using quality mode upscaling at higher resolutions. If you have a 4K monitor, but your GPU isn't up to natively rendering that resolution, upscaling can be a godsend.

Still, there are differences. The best results continue to come from Nvidia's DLSS solutions. If you have an RTX graphics card and are playing a game that supports DLSS, we see little reason not to enable at least quality mode upscaling. At higher resolutions—or lower resolutions on slower RTX GPUs—you'll often get a 30 to 50 percent boost to frame rates. Of course, you can only try DLSS if you have an Nvidia RTX card.

XeSS quality tends to be good as well, and in our experience if you're using an Arc GPU, it often looks slightly better than FSR 2. However, we've also seen games where FSR 2 outperforms XeSS at the same settings, so there's the option to pick slightly higher fps over slightly higher upscaling fidelity. For non-Arc GPUs, XeSS running in DP4a mode is harder to nail down—especially since a lot of XeSS games are still using the older version 1.0 algorithm that looks noticeably worse in DP4a mode. Our testing also shows DP4a mode on AMD and Nvidia GPUs is often more taxing than FSR 2 upscaling on those same GPUs, and XeSS 1.0 typically looks worse than FSR 2.

FRAME GENERATION SHENANIGANS

Some people hate the idea of 'fake pixels' and upscaling on principle. If you're in that camp then so-called frame generation is an even bigger concern. The issue is that it introduces latency and doesn't improve the feel of a game. That's because the algorithms need to interpolate between already rendered frames.

Both DLSS 3 and FSR 3 frame generation work in a similar fashion in that the game renders two frames in the normal fashion—with upscaling if that's enabled—and then passes those frames to an interpolation algorithm that generates an intermediate frame. There are differences after that point.

With Nvidia RTX 40-series GPUs, an AI-trained algorithm runs on a dedicated Optical Flow Accelerator that performs INT8 matrix calculations at a speed of 305 teraops. Nvidia hasn't provided a ton of detail on the specifics, but we do know that the OFA takes the complete frames, including UI elements, and attempts to infer things like motion vectors and other aspects to generate a frame.

AMD's FSR 3 takes the frames without UI elements, which seems safer. There's no machine learning involved, instead using a hand-coded algorithm to interpolate an intermediate frame. The UI elements get overlaid in the usual way.

The generated frame in either case comes without any new user input, and the second of the rendered frames needs to be delayed for this to work. In practice, you end up with an extra two frames of latency (at the generated fps rate). There's also other overhead, so what you typically get is about a 50–70 percent boost in frames to monitor rates. That also means a drop in the base frame rate, leading to a drop in the user input sampling rate.

The quality of the generated frames often depends on the amount of movement in the game. Walking forward in the game world without any fast turns means any two consecutive rendered frames should be similar. Interpolating a frame between the two is thus relatively easy and generally looks good. Fast movement or a camera swap on the other hand can confuse things.

DLSS 3 attempts to interpolate a new frame, regardless of how different the two rendered frames might be. In a worst-case scenario, take *Flight Simulator,* where you change the camera view. Practically everything changes between the two frames, leaving no basis for interpolation. However, DLSS tries to do its job, and often makes up a garbage frame that flashes into view. FSR 3 appears to be more intelligent, detecting major changes and repeating the first frame, effectively turning off for a frame.

In practice, the look and feel of frame generation depends on the game. Most games that use DLSS 3 need a base frame rate of 40 fps or more to feel okay, meaning a generated frame rate of 80 fps or more is viable. The games would still feel like they're running at 40 fps, but that's 'fast enough'—for most gamers, at least. A generated frame rate of only 40–50 fps, meanwhile, ends up feeling very sluggish in our experience.

Frame generation is most beneficial when you're already running at relatively high frame rates—and conversely becomes a lot less useful on lower tier hardware like an RTX 4060 or RX 7600 card. 4K upscaled output running at 30–40 fps after frame generation means the game would feel like it's running at 15–20 fps, and that's not something we consider playable.

WELCOME TO YOUR NEURALLY RENDERED FUTURE

Technologies like DLSS, FSR, and XeSS fundamentally alter a lot of the normal assumptions we make about what settings to use when playing games. Some will tell you playing at anything other than native resolution is wrong, but the more we've poked at the

©GEARBOX/THQ NORDIC, TECHLAND

Dying Light 2 Stay Human

Naraka: Bladepoint

various options, the less weight we're willing to give such claims. The reality is that graphics have always been about 'cheating' and image quality 'hacks.' Is a good upscaling solution any different?

Consoles have been using dynamic upscaling algorithms for years, and complaints from PC gamers about how consoles are inferior feel more like sour grapes than anything concrete. Sure, a PC with the latest AMD or Intel CPU and an RTX 4090 offers more computational and graphical power than an Xbox Series X or PlayStation 5. The GPU alone also costs, at present, about four times as much as an entire gaming console. To each their own.

We tend to be open minded about what settings, resolutions, and other technologies we enable when playing games. Some games use fast and dirty techniques that clearly reduce image quality. Others go all-in on ray tracing and whizbang graphics effects. Ultimately, does the tech behind a game even matter if you're enjoying the game?

Upscaling, frame generation, ray reconstruction, and future neural rendering techniques are ultimately tools for game developers. Some games may misuse those tools to try to hide other problems. Others will choose not to implement such tools because they're deemed unnecessary. All we need to do is look at some of the advances we're seeing around the world thanks to AI and deep learning to know that these things aren't going away.

What will come next? Nvidia is already working on neural texture compression that could potentially reduce texture memory use by 75 percent. Intel wants to extrapolate new frames. Other companies are working to meld ChatGPT-like devices with text to speech and speech to text. AI level design, graphics, stories, and even game code are all down the pipeline.

Some games and techniques will come up short. Others will push boundaries. We're still excited to see what the future holds—and I'm still waiting to be able to jack into cyberspace. ⏻

One notable game that supports XeSS is *Hogwarts Legacy*.

FRAME EXTRAPOLATION

XeSS doesn't offer frame generation, but Intel says it's working on a different algorithm: frame extrapolation. Rather than interpolating between two generated frames, frame extrapolation would attempt to generate—via AI—the next frame after the current frame.

It's ambitious, but making it work will require some clever algorithms. Still, we can see the benefits. There would be no added latency, and the creation of the generated frame could run concurrently with the rendering of the next frame.

How would extrapolation work? In its Siggraph paper, 'ExtraSS: A Framework for Joint Spatial Super Sampling and Frame Extrapolation', researchers discuss methods of 'warping' to generate a future frame from the current frame. Early results seem promising, though it could be a year or more before a public version is available.

Besides making extrapolation viable, Intel faces an uphill battle in getting developers to use it. It's the smallest, in terms of market share, of the three GPU manufacturers. Making an Arc-exclusive feature would cater to a minuscule part of the market, so perhaps Intel will elect for an open-source approach.

Given the current Arc A-series graphics cards are already on the lower end of the performance spectrum, this could be something for next-gen Battlemage GPUs. Those will presumably be more capable for both graphics and AI workloads, and the added horsepower could make extrapolation a reality. We can't help but wonder if Nvidia might also be working on it as part of a future DLSS upgrade.

© NETEASE GAMES, PORTKEY GAMES LABEL

magazinesdirect.com

Over 100 brands to choose from

Computer active
GUITAR WORLD
APOLLO MISSIONS
T3
retro GAMER
PC GAMER

GIFT IDEA! Gifts that keep on giving

3 GREAT REASONS TO SHOP WITH US

- From tech and gaming titles to fashion and celebrity magazines there's something for everyone.
- Immerse yourself in our specialist one-off publications about your favorite hobby or interest.
- Exclusive Offer - Get 15% Off Guides and Specials - **MDMAG15**

VISIT OUR ONLINE STORE AT
www.magazinesdirect.com

✓ Convenient home delivery ✓ Save on shop price ✓ Never miss an issue ✓ 1.4m subscriptions sold

MAXIMUMPC

PREMIUM DIGITAL SUBSCRIPTION

As a premium subscriber you instantly get access to **100+ back issues**

Z790 MOBOS REVIEWED PG.74, 84
WINDOWS COPILOT FULL GUIDE PG.58
HOW TO BAN ADS PG.62
MAXIMUMPC
MINIMUM BS • JANUARY 2024 • www.maximumpc.com
BUILD A TOUCH SCREEN GAMING PC
DO IT YOURSELF! STEP-BY-STEP GUIDE PG.16
Core i5-14600K + RTX 4080
4K Gaming Powerhouse
14th Gen i5 vs i9 Compared
LIVESTREAM LIKE A PRO PG.50
WIN 10: STICK OR SWITCH? PG.42
RX 7700 XT REVIEW PG.74
BUILD TOP-E AMD P
RX 7900 XTX + Ryzen 9 7900X
Maxed out AMD for just $3,000
VRAM: HOW MUCH DO YOU NEED? PG.30
WEB BROWSER FACE-OFF PG.50
RASPBERRY Pi 5 REVIEWED PG.82
4TH GEN MING PC
STORE IMAGES FOR FREE! PG.64 & 70

$34.95* for one year/ 13 issues

www.magazinesdirect.com/MPP

*Price in US $. Offer is valid on US orders only, visit us at www.magazinesdirect.com or call 1-844-779-2822 for other subscription options.
Offer valid until February 30 2024.

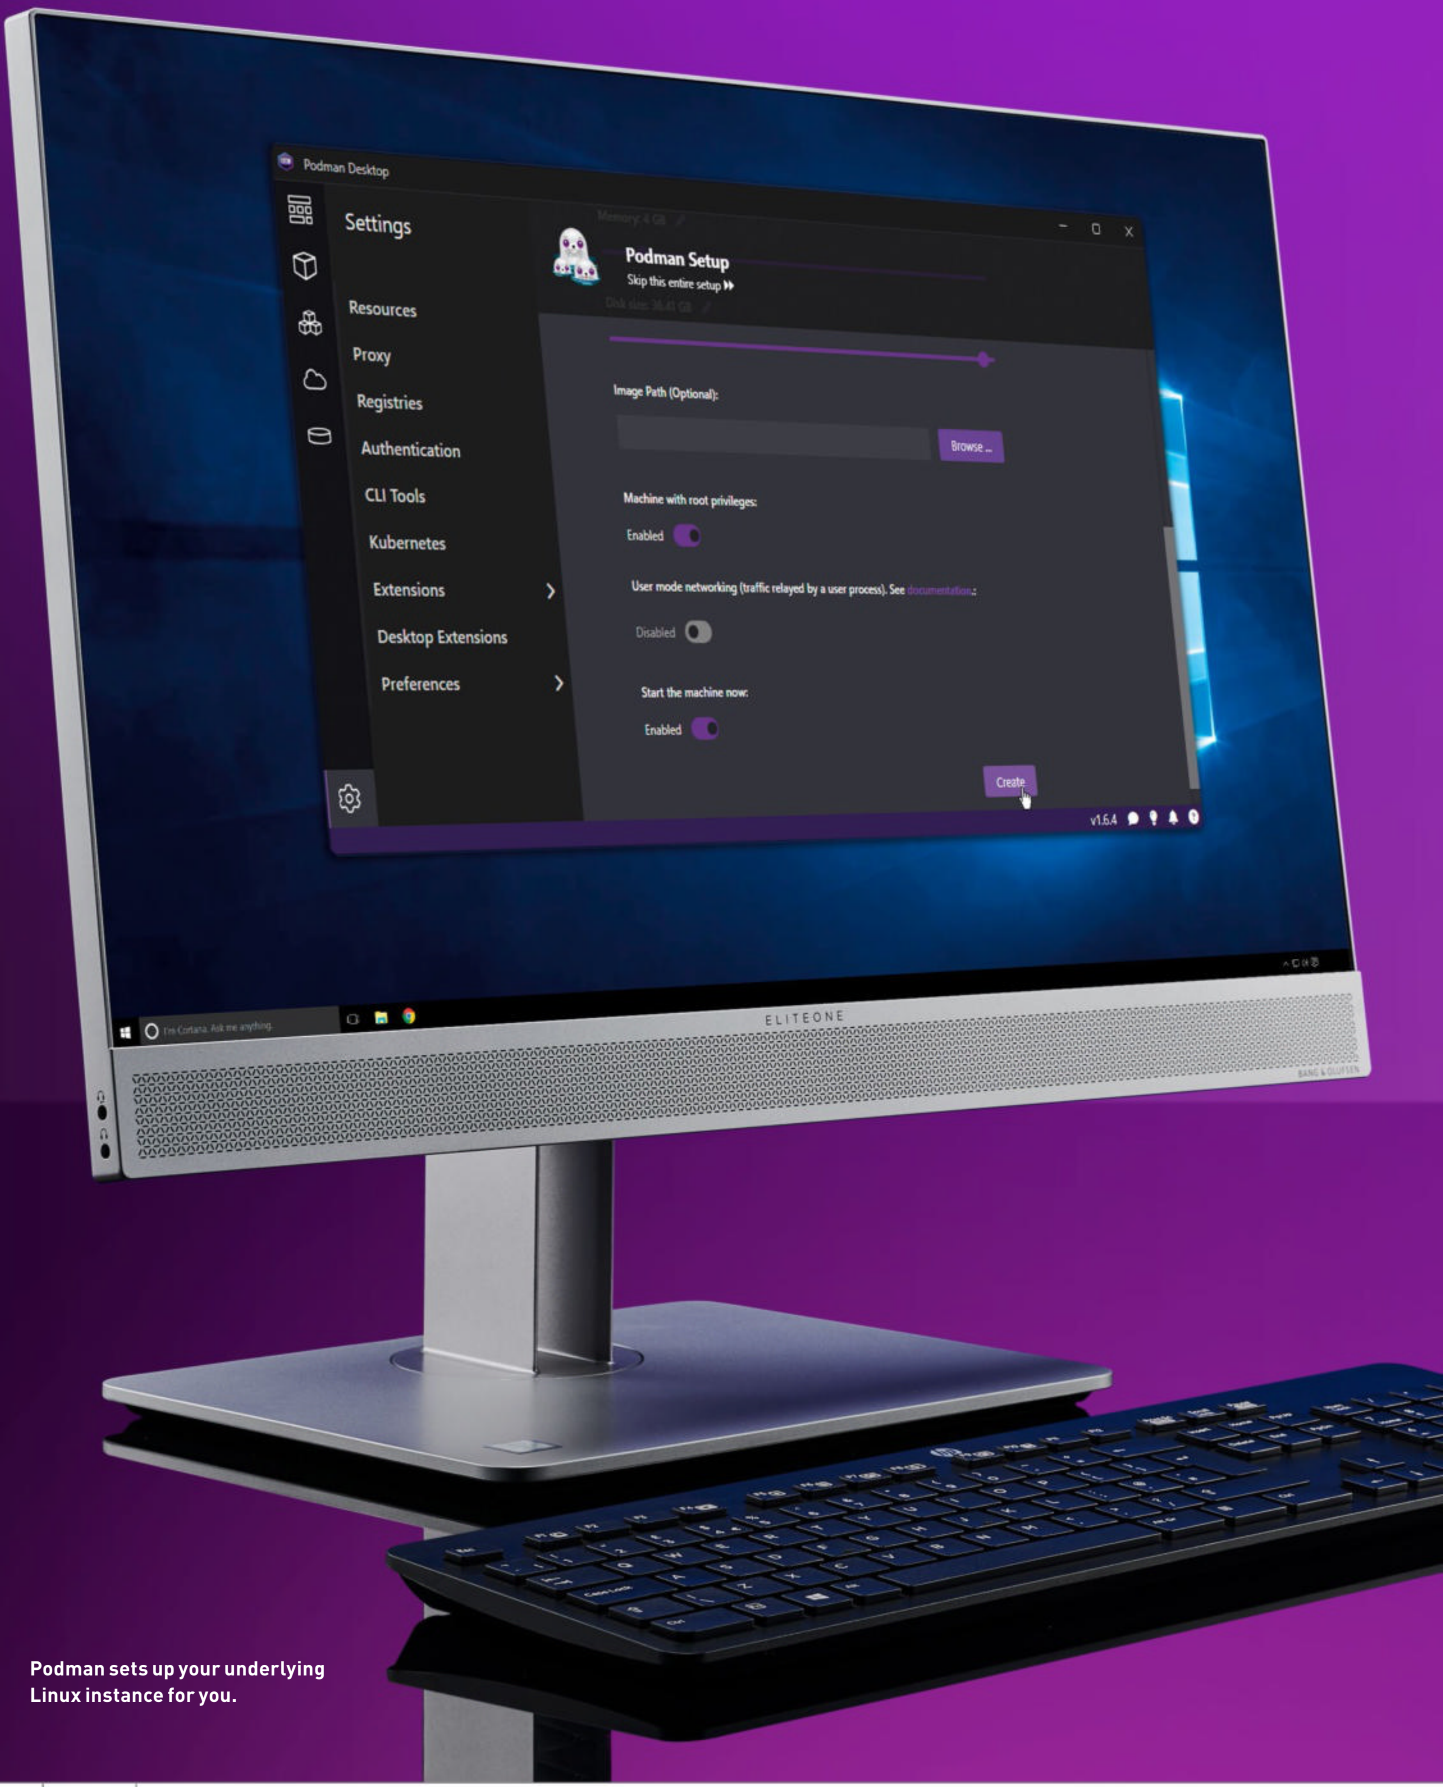

Podman Desktop
Settings
Resources
Proxy
Registries
Authentication
CLI Tools
Kubernetes
Extensions
Desktop Extensions
Preferences
Podman Setup
Skip this entire setup
Browse ...
Create
v1.6.4
ELITEONE

Image Path (Optional):

Machine with root privileges:
Enabled

User mode networking (traffic relayed by a user process). See documentation.:
Disabled

Start the machine now:
Enabled

Podman sets up your underlying Linux instance for you.

THE BEGINNER'S GUIDE TO PODMAN

Discover a new, more secure way to run containers with *Nick Peers*

WE'RE HUGE FANS of containers. They're a great way to run self-hosted services on your network, and over the past few years we've covered them extensively, from Nextcloud (cloud storage, chat and web office suite) and Immich (photo storage), to Vaultwarden (password managing) and Motion (home surveillance).

In the past, we've focused on Docker as the tool for running these, but there's a new pretender on the block. What Podman lacks in finesse it makes up for by offering superior performance and security with 'rootless' containers—see our head-to-head in the December 2023 issue.

Having come out in Podman's corner, we thought it only fair to give you a lowdown of the basics of using it on your Windows PC. In this feature, we'll introduce you to the concept of containers, and how to set up and run self-hosted services using a mixture of the user-friendly Podman Desktop tool and command line engine.

We'll also reveal how to deal with some of the idiosyncrasies of rootless containers, helping you get to grips with core features like autostarting containers with Windows and implementing a reverse proxy to give you access to your services from outside your home.

© PODMAN

LET'S START BY DEFINING what a 'container' is. Think of it as a stripped-back virtual machine—while a tool like VirtualBox emulates an entire computer's hardware and its underlying operating system, a container contains the bare minimum needed to run whatever application is required of it. Containers are Linux-based, so when you run them in Windows, you'll need the Windows Subsystem for Linux (WSL) installed. From here, Podman will create a 'machine', a minimal version of Linux (Fedora, if you're interested), inside which your containers can run with just the specific files and dependencies they need. This means a container takes up tens of megabytes of space instead of gigabytes, and that the container consumes far fewer resources than a virtual machine, and any applications run inside it.

Although they share access to the underlying kernel, containers are designed to run in complete isolation to each other. In the case of Docker, a single daemon runs on top of these, and containers are given root access by default; Podman's selling points are greater security through rootless containers, plus greater stability and performance by eliminating the daemon process.

Podman requires Windows 10 or 11, with hardware virtualization enabled. To confirm this, right-click the taskbar and choose Task Manager, then switch to the Performance tab and select CPU. Check the Virtualization setting—if it's Enabled, then it's switched on; if it's Disabled, you'll need to reboot into your UEFI/BIOS settings to enable hardware virtualization. Consult your PC or motherboard manual for where to find it (look for a setting related to Intel VT, Intel VT-x, or AMD-V, depending on your CPU).

Install Podman on Windows

From here, installation is a simple affair. Head over to https://podman.io and click Download to reveal two options: Podman Desktop and Podman CLI. For full functionality and a user-friendly management interface, choose Podman Desktop. Once installed, the Podman Desktop window will appear, and you'll be told that the underlying Podman engine requires installation.

Click 'Set up', followed by Next, to install Podman. A series of checks will be run, and you'll be invited to install the latest build (4.8.3 at time of writing). Click Yes, make sure 'Install WSL if not present' is checked, and click Install. If needed, the installer will install WSL—a reboot may be required—after which, click Close.

Leave Autostart enabled, then click Next. You'll be invited to create a Podman machine—click Next to bring up the wizard. Scroll down, leaving the default settings alone, then click Create to create your virtual machine (a minimal version of Fedora, inside which all containers will run). Click Next when done.

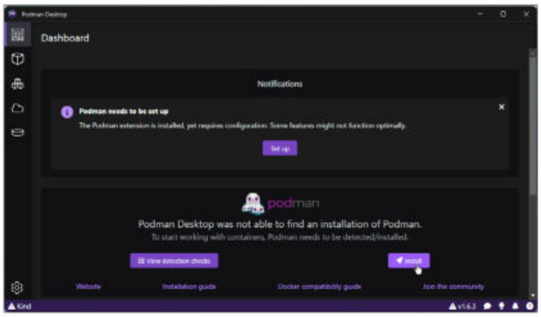

Podman needs to be set up

The Podman extension is installed, yet requires configuration. Some features might not function optimally.

Podman Desktop was not able to find an installation of Podman.

To start working with containers, Podman needs to be detected/installed.

Once the desktop is in place, install the underlying engine.

Your first container

Podman can be run directly through the command line or via Podman Desktop. The latter is more user-friendly, and is useful for learning how Podman creates and runs containers, so we'll start there. We're going to install a perennial favorite as our first container: Vaultwarden, a lightweight (and unofficial) self-hosted implementation of open-source password manager Bitwarden.

As with all containers, visit its web page first to see how it's put together—a good place to start looking is the official Docker hub, where most popular containers (including Vaultwarden—see https://hub.docker.com/r/vaultwarden/server) reside. In the case of Vaultwarden, we can see that it's created by using two Docker commands:

```
docker pull vaultwarden/server:latest
docker run -d --name vaultwarden
-v /vw-data/:/data/ -p 9000:80
vaultwarden/server:latest
```

MIGRATING FROM DOCKER

If you've been happily setting up Docker containers following our guides over the past few years, you'll be pleased to learn that moving them across to Podman is simple—if you point your Podman container to the same physical folders you used for Docker, the switch should be seamless.

Before going further, bookmark https://docs.podman.io/en/latest/index.html—this provides a handy reference to all supported Podman commands. Most work in virtually the same way as those issued using the 'docker' command—this is by design, as Podman's developers are keen to make transitioning from Docker as smooth as possible.

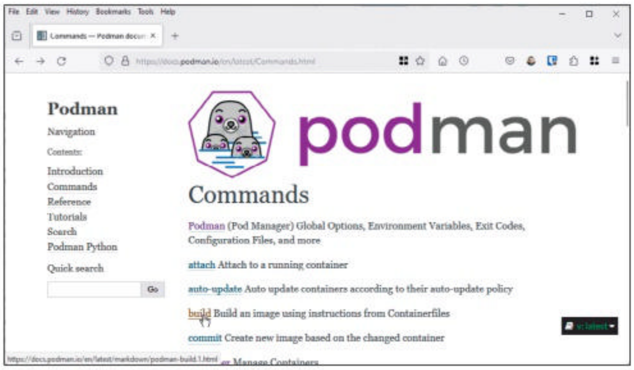
File Edit View History Bookmarks Tools Help
Commands — Podman docu
https://docs.podman.io/en/latest/Commands.html
Podman
Navigation
Contents:
Introduction
Commands
Reference
Tutorials
Search
Podman Python
Quick search
Go
podman
v: latest
https://docs.podman.io/en/latest/markdown/podman-build.1.html

Commands

Podman (Pod Manager) Global Options, Environment Variables, Exit Codes, Configuration Files, and more

attach Attach to a running container

auto-update Auto update containers according to their auto-update policy

build Build an image using instructions from Containerfiles

commit Create new image based on the changed container

Podman commands are designed to closely follow Docker's.

That's the theory, but we've already seen problems with networking and automatically restarting containers that are unique to Podman in Windows. This is down to both the daemonless approach and Podman's support for rootless containers.

Another consequence of rootless containers is that you can't use any ports below 1,024 for mapping—this is an arbitrary figure chosen by Podman that is likely to be relaxed, but you can map these to higher ports (such as port 9000 for Vaultwarden). If that's an issue, you'll find a workaround involving podman-privileged-ports.conf at www.smarthomebeginner.com/docker-to-podman-migration-guide/#Networks—you'll need to 'podman machine ssh' into your WSL instance to issue it.

If you've used Docker Compose or Portainer's Stacks feature to group related containers together (for example, to create a Nextcloud AIO instance), then the equivalent in Podman is Pods using Kubernetes YAML files.

© PODMAN

The first command reveals the name of the 'image' used to create the Vaultwarden container, so our first task is to pull this image into Podman. This is done by clicking the cloud icon on the left of the Podman Desktop interface to navigate to the Images section. Click the Pull button in the top right corner, where you'll be invited to input the image name. Type vaultwarden/server here and click 'Pull image'—you'll see the image downloaded in six parts—click Done when it's finished.

You'll now see the Vaultwarden image appear in the list as docker.io/vaultwarden/server. The docker.io portion refers to the registry that Podman pulled Vaultwarden from—in this case, the official Docker Hub. If you click the Settings button in the bottom left corner followed by Registries, you'll see Docker Hub is one of four pre-configured registries supported, the others being Red Hat Quay, GitHub, and Google Container Registry. You can connect an existing account to any of these using the Configure button next to the relevant entry, but they should all work without one.

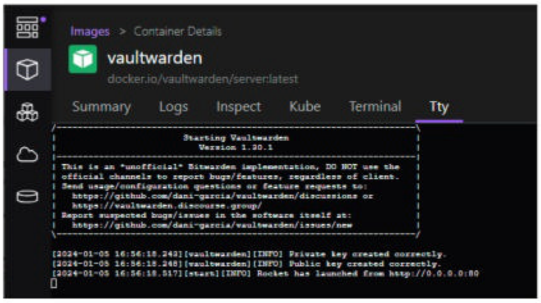
Images > Container Details
vaultwarden
docker.io/vaultwarden/server:latest
Summary
Logs
Inspect
Kube
Terminal
Tty

```
/--------------------------------------------------------------------\
|                        Starting Vaultwarden                        |
|                           Version 1.30.1                           |
|--------------------------------------------------------------------|
| This is an *unofficial* Bitwarden implementation, DO NOT use the   |
| official channels to report bugs/features, regardless of client.  |
| Send usage/configuration questions or feature requests to:         |
|   https://github.com/dani-garcia/vaultwarden/discussions or        |
|   https://vaultwarden.discourse.group/                             |
| Report suspected bugs/issues in the software itself at:            |
|   https://github.com/dani-garcia/vaultwarden/issues/new            |
\--------------------------------------------------------------------/

[2024-01-05 16:56:18.243][vaultwarden][INFO] Private key created correctly.
[2024-01-05 16:56:18.248][vaultwarden][INFO] Public key created correctly.
[2024-01-05 16:56:18.517][start][INFO] Rocket has launched from http://0.0.0.0:80
```

Once set up, use the Tty tab to confirm that it's working.

Set up Vaultwarden

We're now ready to 'run' Vaultwarden as a container. This is done by clicking the play button to the right of its entry to bring up the Run Image wizard. This is split into four sections, but for the purposes of launching Vaultwarden for the first time, we only need to focus on the Basic tab. Everything you need to replicate from the 'docker run' command can be found here.

Start by populating the 'Container name:' tab with 'vaultwarden'. Leave the next two fields ('Entrypoint' and 'Command') alone, and focus on the Volumes section. Volumes reveals an important part of how containers work: the container itself is immutable, which means that every time you shut it down and restart it, the container (along with its contents) is destroyed and recreated. Volumes allow you to store your settings in a physical folder on your PC, where data can be amended. These folders are then 'mapped' to folders inside the container, which ensures that your settings and any other relevant data survive the container's destruction.

Vaultwarden stores all your configuration data—including your encrypted password database file—inside such a volume. First, click the folder icon inside the 'Path on the host' field to select a folder on your hard drive to keep this data. We recommend setting up a dedicated Containers folder, inside which you create a Vaultwarden folder for this purpose.

Next, click inside the 'Path inside the container' field to input the path you're mapping your folder to inside the container—in the case of Vaultwarden, it's /data/. Vaultwarden only requires a single folder mapped, but other containers may require more, in which case you'd just click the + button to the right of this entry to add additional entries.

Mapping network ports

We come to the 'Port mapping' section. Containers communicate with the outside world through ports, and because there's a danger that some containers may wish to communicate on the same port (port 80 is popular), you can map different ports on your PC to ports inside the container to allow communication to flow freely.

Unlike volumes, container images usually pre-fill the ports you'll need—in the case of Vaultwarden, both ports 3012 and 80 are required, and you'll see that while it's a straight mapping of port 3012 to 3012, port 80 is remapped to 9000. This is because Podman makes all ports from 1 to 1023 off-limits to rootless containers for security purposes, although it's likely to relax this rule in a later release. For now, however, you can leave the ports as they are, or specify your own alternatives, if you prefer—in most cases, they can be left alone.

COMMAND-LINE USAGE

Recreating containers using Podman Desktop is a faff. The good news is that if you learn how to use Podman from the command line, you can recreate containers in seconds rather than minutes. Even better, Podman can be invoked from any shell, including both Powershell and the Command Prompt. Simply right-click the Start button, choose Terminal (Windows 11) or Command Prompt (Windows 10), and you're good to go.

Containers are created (and recreated) using the 'podman run' command. Rather than type these in manually each time, we recommend setting up text files, inside which you store commands like so:

```
podman run -d --name vaultwarden -v C:/users/username/containers/vaultwarden:/data/ -e SIGNUPS_ALLOWED=false -p 9000:80 --restart always docker.io/vaultwarden/server:latest
```

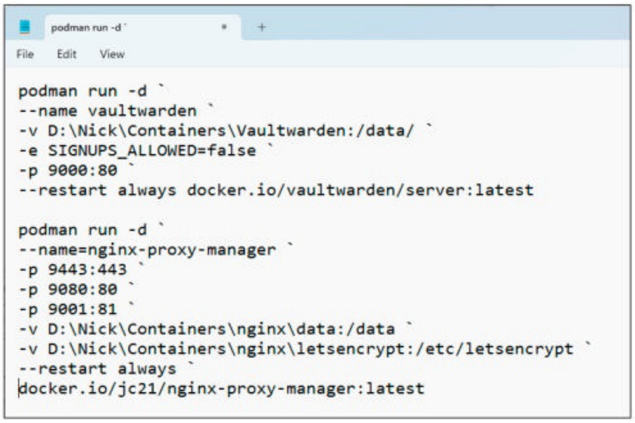

```
podman run -d `
--name vaultwarden `
-v D:\Nick\Containers\Vaultwarden:/data/ `
-e SIGNUPS_ALLOWED=false `
-p 9000:80 `
--restart always docker.io/vaultwarden/server:latest

podman run -d `
--name=nginx-proxy-manager `
-p 9443:443 `
-p 9080:80 `
-p 9001:81 `
-v D:\Nick\Containers\nginx\data:/data `
-v D:\Nick\Containers\nginx\letsencrypt:/etc/letsencrypt `
--restart always `
docker.io/jc21/nginx-proxy-manager:latest
```

Store your Podman run commands in a text file for easy access.

If this is difficult to follow, you can break this into multiple lines using the ` symbol (see the screenshot).

You can then simply amend the command as required before copying it to the clipboard, switching to your shell window, and pressing Shift + Insert to paste it in. Favor Powershell over Command Prompt, and the command is color coded, helping you to quickly verify it before hitting Enter.

Before recreating a container, you need to stop and remove it first—if Podman Desktop isn't available, issue the following two commands:

```
podman stop vaultwarden
podman rm vaultwarden
```

Below this is a section called 'Environment variables'—this can be used to define all manner of custom settings. Vaultwarden supports an environment variable called SIGNUPS_ALLOWED—this allows you to prevent others from signing up for their own password accounts by setting it to 'false'. If it's not defined, signups are allowed, which you'll need to set up your own account, so leave Environment variables blank for now.

You now have everything you need to get your Vaultwarden container up and running. Click 'Start Container', and you'll be taken to the Tty tab of your Vaultwarden's container, where you'll see a series of status messages appear that should hopefully confirm the container is up and running—in the case of Vaultwarden, you should see a 'Rocket has launched from http://0.0.0.0:80' message.

This indicates Vaultwarden is listening internally on port 80 (which corresponds to port 9000 on your PC), so open your browser and go to http://localhost:9000 where you should find Vaultwarden is ready and waiting for you to set it up by clicking 'Create account'. Do so, and you will find yourself at the web vault screen.

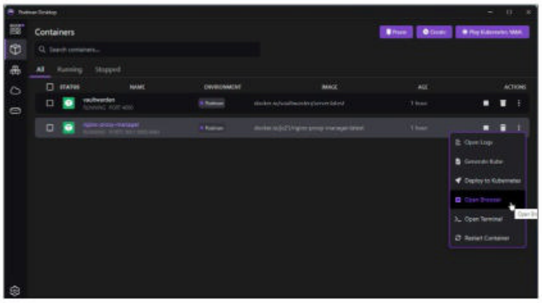

View and manage running containers from here.

Network fix

What you can't yet do is access Vaultwarden from any other device on your network. This is a problem peculiar to the WSL implementation Windows uses (those running Podman on a native Linux install won't have this issue), and is fixed by binding specific ports to the underlying Linux machine's IP address.

To discover what this is, open an elevated Powershell window as an administrator—Windows 11 users can do this by right-clicking the Start button and choosing Terminal (Admin). Now type the following command:

```
wsl hostname -I
```

(If more than one entry appears, the first one—typically 176.23.x.y—is the correct one.)

Armed with this information, you can 'bind' the port using another Powershell command. The following example 'binds' port 9000 to a WSL instance with an IP address of 172.23.247.76:

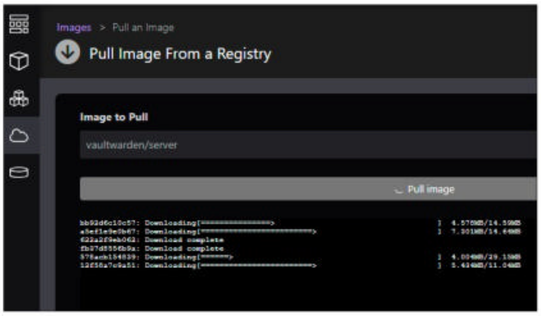

Before a container is created, its image must be pulled.

```
netsh interface portproxy add v4tov4
listenaddress=0.0.0.0 listenport=9000
connectaddress=172.23.247.76
connectport=9000
```

When the command is issued, you should be able to access your Vaultwarden instance from another device on your network using your PC's IP address—for example, http://192.168.0.2:9000. But while Windows now does its best to maintain that IP address across reboots, at some point the IP address of your WSL instance may change, and the fix will stop working. When this happens, issue the following command in an elevated Powershell window:

```
netsh interface portproxy reset
```

Then, repeat the 'wsl hostname -I' command to get the new IP address and reissue the 'netsh interface portproxy add' command to point to its new address.

As you add more containers, you will find yourself opening more and more ports. You can review what these are

TROUBLESHOOT PROBLEMS

Podman is a relatively new product, and is still undergoing rapid development—both the desktop app and underlying engine were updated as we were writing this feature. Because it's a child of Red Hat Linux, development is focused primarily on Linux (and Fedora/Red Hat Linux), which shows in how the project is developed for Windows.

The need to use an underlying WSL instance creates unique problems—we've covered the two major issues (automatic container restarting and remote network access) here, but you're likely to run into other problems too. Remember, you can destroy and recreate containers, but if you mangle your underlying WSL instance, you can destroy and recreate that, too: navigate to 'Settings > Resources' in Podman Desktop and click the stop button underneath the Podman machine, before clicking the trash to remove it. Once done, click 'Create new…' to set up a new instance. If you go down this route, you'll need to reapply all network fixes and any commands issued inside the instance via 'podman machine ssh'.

When it comes to troubleshooting, start by visiting the troubleshooting page in Podman Desktop (see https://podman-desktop.io/docs/troubleshooting for details), then expand out to the web. There's Podman's Github page (https://github.com/containers/podman/issues), but in most cases, a web search (don't forget the 'Windows' keyword alongside Podman) will lead you to solutions or workarounds. If all else fails, why not tap into the Doctor's expertise by emailing your setup and problems to doctor@maximumpc.com?

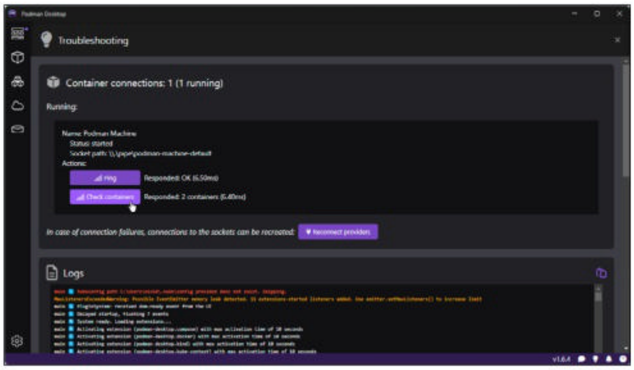

Name: Podman Machine
Status: started
Socket path: \\.\pipe\podman-machine-default
Actions:
Responded: OK (6.50ms)
Responded: 2 containers (6.40ms)

In case of connection failures, connections to the sockets can be recreated:

Podman Desktop has a dedicated troubleshooting section.

© PODMAN

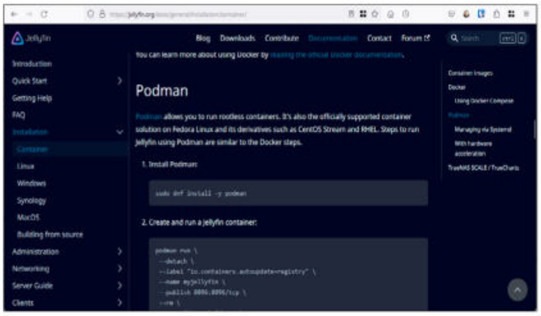

You can learn more about using Docker by reading the official Docker documentation.

Podman

Podman allows you to run rootless containers. It's also the officially supported container solution on Fedora Linux and its derivatives such as CentOS Stream and RHEL. Steps to run Jellyfin using Podman are similar to the Docker steps.

1. Install Podman:

```
sudo dnf install -y podman
```

2. Create and run a Jellyfin container:

```
podman run \
 --detach \
 --label "io.containers.autoupdate=registry" \
 --name myjellyfin \
 --publish 8096:8096/tcp \
 --rm \
```

Look for Podman-specific instructions where they exist.

using the 'netsh interface portproxy show all' command, making it easier to remap them when required. We expect Podman to better handle remote networking connections in a future release, but for now, this workaround will do the job.

Recreating containers

Unlike with Docker, you can close Podman completely, and your container will run independently in the background. Reboot your PC, however, and you will see that it's stopped. You can relaunch it manually by opening Podman Desktop—wait a few minutes, and you will see a series of status messages hopefully resolve themselves as 'Podman is running'. You can switch to Containers and click the play button next to Vaultwarden, but that's a faff.

Instead, you can configure containers to automatically restart whenever they've been stopped, so after rebooting your PC, Vaultwarden will automatically start back up in the background.

To recreate the container with these new settings, delete the current container by clicking the trash icon next to its entry under Containers. Once deleted, switch to Images, and click the play button next to the Vaultwarden image. The Run Image wizard will reappear—fill it in the Basic tab as before, but as you've set up your Vaultwarden account, now is a good time to input SIGNUPS_ALLOWED into the Name field under Environment Variables, followed by false into the 'Value (leave blank for empty)' field, and click the + button.

Now, switch to the Advanced tab, and you'll see a Restart Policy section with a Policy Name drop-down currently set to 'No restart'. Click this and change it to 'Always restart'—in theory, this should restart the container after every reboot. We say 'in theory', because by default, Podman doesn't restart any containers after you restart Windows.

The one-time fix involves creating a service inside your WSL instance to automatically restart rootless containers when the instance restarts alongside Windows. To do this, you need to log into the underlying Linux-based distro to issue some commands, again from a Powershell window. This ensures that any containers set to automatically restart will do so when Podman is started. To enter the instance, type the following:

```
podman machine ssh
```

When the '[root@PC-NAME ~] #' prompt appears, type the following:

```
cp /lib/systemd/system/podman-restart.service /home/user/.config/systemd/user/
systemctl --user daemon-reload
systemctl --user enable --now podman.socket
systemctl --user enable --now podman-restart
loginctl enable-linger $USER
```

The final step is to ensure Podman starts with Windows—go to Settings > Preferences in Podman Desktop, and make sure that 'Start on login' and 'Autostart Podman engine' are enabled.

Beyond the basics

You can experiment with creating other containers, too. When you've found one, Google its name and 'podman' to see if instructions exist (Jellyfin has a Podman section at https://jellyfin.org/docs/general/installation/container). If not, try replicating the Docker instructions.

While you can access your web vault on your machine via localhost, when you try to log into it through another device or your PC's IP address, you'll be told the browser requires HTTPS to access it. The Vaultwarden wiki (https://github.com/dani-garcia/vaultwarden/wiki) reveals one of two options—the easiest is to set up a reverse proxy, and the final box reveals how to do that not just for Vaultwarden, but for any services that require access from outside your home network. ⏻

SET UP A REVERSE PROXY

Many services you set up work best when they can be accessed from outside your home network. We've covered reverse proxies before—they basically make it possible to redirect traffic from outside your network to the appropriate service. Our favorite proxy is the user-friendly Nginx Proxy Manager, which comes in container form.

For this to work, you need to configure both Vaultwarden and the Nginx Proxy Manager to be able to communicate with each other. This can be done by running both containers in a 'pod'. Assuming your Vaultwarden instance is up and running, the next step is to set up Nginx Proxy Manager.

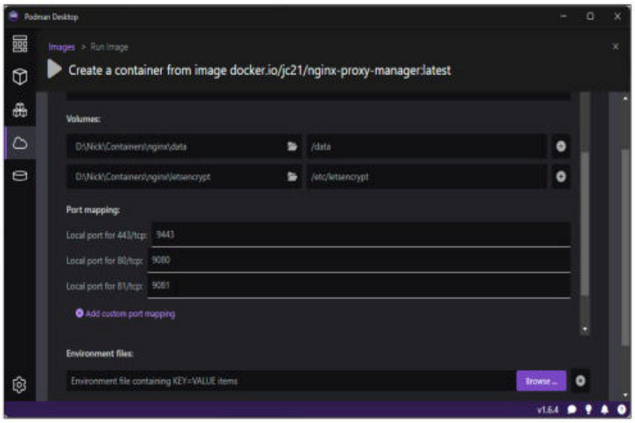

Create a container from image docker.io/jc21/nginx-proxy-manager:latest

Volumes:
D:\Nick\Containers\nginx\data → /data
D:\Nick\Containers\nginx\letsencrypt → /etc/letsencrypt

Port mapping:
Local port for 443/tcp: 9443
Local port for 80/tcp: 9080
Local port for 81/tcp: 9081

Environment files:
Environment file containing KEY=VALUE items

Nginx Proxy Manager is our favored reverse proxy.

First, navigate to Images in Podman Desktop and click Pull. Type the following into the Image Name field: docker.io/jc21/nginx-proxy-manager:latest. Once the image is downloaded, click Play and set it up as follows:

Basic tab: name it nginx-proxy-manager, and set up two volumes, one pointing to /data, the other to /etc/letsencrypt. Remap ports 443, 80 and 81 to ports higher than 1024. Advanced tab: set Policy Name to 'Always restart'. Click Start Container.

Go to http://localhost:9081/ to verify Nginx Proxy Manager is up and running. Head to https://linuxformat.com/archives?listpdfs=1 where you can download our sister title Linux Format's four-page guide to configuring and setting up Nginx Proxy Manager for free, allowing you to use Vaultwarden and other container services from outside your local network.

RASPBERRY

Your next desktop PC for $60

The lightweight, British-made Pi is perfect for everyday desktop duties. *Nik Rawlinson* explains how to get started

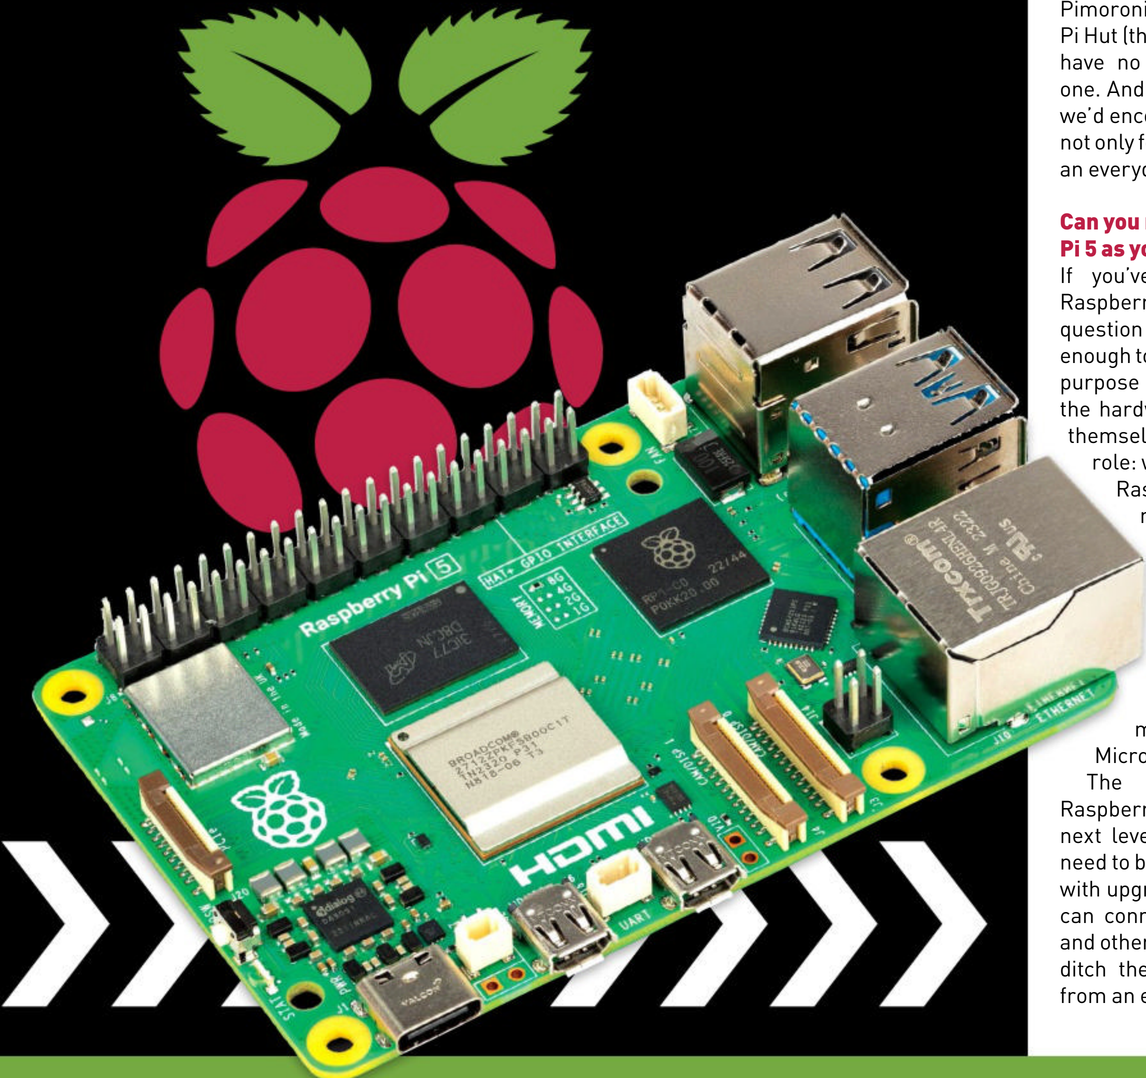

© RASPBERRY PI FOUNDATION

THE NEW RASPBERRY PI is here, and it's better than ever, with a big boost in processor power and enhanced peripheral support. What's more, the price is still absurdly low. It's hardly surprising that when the first shipment of Raspberry Pi 5 boards arrived in October 2023, they sold out almost immediately—although that might also have something to do with the post-pandemic chip shortage, which for the past few years has made most Raspberry Pi models almost impossible to buy.

Happily, more Raspberry Pi 5 boards are already hitting the shelves via suppliers such as Pimoroni (pimoroni.com) and The Pi Hut (thepihut.com), so you should have no difficulty getting hold of one. And perhaps for the first time, we'd encourage you to consider one not only for hobbyist projects, but as an everyday desktop computer.

Can you really use the Raspberry Pi 5 as your desktop PC?

If you've tried using an older Raspberry Pi board, you might question whether it's fast or flexible enough to use as a regular, general-purpose PC. But recent versions of the hardware have already proven themselves viable for a desktop role: when the pandemic struck, Raspberry Pi said it saw a rapid increase in the use of Raspberry Pi 4 for home working and studying, and in November 2020 it unveiled the Raspberry Pi 400, with a compact keyboard inspired by classic single-box machines, such as the BBC Micro and ZX Spectrum.

The new hardware in the Raspberry Pi 5 takes things to the next level, with all the power you need to be productive. It also comes with upgraded USB support, so you can connect fast external storage and other peripherals: you can even ditch the microSD card and boot from an external SSD.

Software support keeps getting better, too. The official operating system goes from strength to strength, while free software, such as LibreOffice allows you to collaborate effortlessly with people using Microsoft Office. The bundled version of Firefox is now optimized for Raspberry Pi, with particular attention paid to desktop sharing and video-call performance—ideal for working from home. If you're running Ubuntu elsewhere, you have the option of installing the exact same OS on the Raspberry Pi hardware.

Perhaps there will always be particular use cases and programs that work better in Windows or macOS. But the Pi is maturing, and at the same time more applications are evolving into web apps and cloud services. It's never been more enjoyable or practical to use the Raspberry Pi as your primary desktop computing platform.

What's new in the Raspberry Pi 5?

We've already given away the Pi 5's headline feature: its speed. The fastest Raspberry Pi 4 model was built on an ARM Cortex-A72 chip clocked at 1.8GHz, but the new board uses a more advanced Cortex-A76 CPU running at 2.4GHz. That's a huge step up in computing performance; the GPU and system memory are also faster, and some core functions have been moved onto dedicated chips, providing further efficiency gains.

The new Pi provides some new interfaces, too. There are now two four-lane MIPI DSI/CSI connectors where the headphone jack and camera connector sat, and a PCI Express 2 interface where you'd have found the display connector. Wi-Fi 5 and Bluetooth 5 are built in, as before, but there's also now a battery-powered real-time clock module, and—for the first time—a power button built into the board itself. Pressing this once brings up the Raspberry Pi OS shutdown dialog; a second press immediately launches a clean shutdown.

Some other ports have been moved. The gigabit Ethernet and USB ports have swapped places, returning to the positions they occupied on the Pi 3 and earlier boards, while the 3.5mm headphone jack has disappeared, so you'll have to rely on Bluetooth or HDMI for sound output. The USB ports are now all full-size Type-A sockets, with two USB 2 sockets and two supporting USB 3; there's also a sole USB-C port, but that's only there for power. The Pi 5 costs from $25 more than the 2021 Raspberry Pi 4 Model B: you'll pay $60 inc VAT for the 4GB model, or $80 for the 8GB board.

The ideal desktop setup

If you're buying a Pi 5 for desktop use, there aren't many decisions to make. The computer itself currently comes only as a bare board (although we'd love to see a Pi 500 model with an integrated keyboard), in 4GB and 8GB variants. You can't upgrade the RAM after purchase, so we'd recommend the extra headroom of the 8GB option.

As with previous models, you can power the Raspberry Pi 5 from almost any USB-C supply, but the new model adds support for the USB-PD standard, and raises the power limit from 15W to a maximum of 27W, allowing it to power a full set of USB peripherals, such as external hard drives. The official 27W supply can be had for $12.60.

Want to put your desktop Pi in a case? You probably won't be able to reuse an old one thanks to the updated board layout, but Raspberry Pi has revamped its distinctive red and white case to incorporate a heatsink and fan for just $10.48. Alternatively, there's already a good range of compatible third-party cases to choose from: it was also $10.48 for the open-topped Pibow Coupe 5 case (tinyurl.com/54bahb78), which gives easy

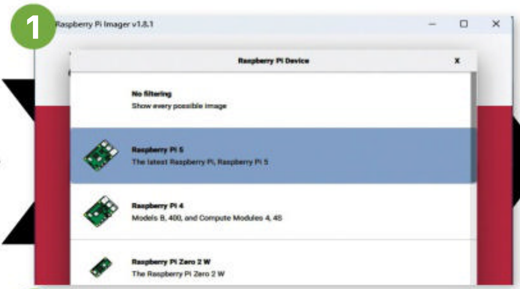

No filtering
Show every possible image

Raspberry Pi 5
The latest Raspberry Pi, Raspberry Pi 5

Raspberry Pi 4
Models B, 400, and Compute Modules 4, 4S

Raspberry Pi Zero 2 W
The Raspberry Pi Zero 2 W

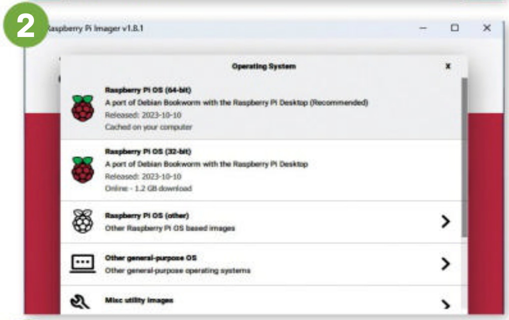

Raspberry Pi OS (64-bit)
A port of Debian Bookworm with the Raspberry Pi Desktop (Recommended)
Released: 2023-10-10
Cached on your computer

Raspberry Pi OS (32-bit)
A port of Debian Bookworm with the Raspberry Pi Desktop
Released: 2023-10-10
Online - 1.2 GB download

Raspberry Pi OS (other)
Other Raspberry Pi OS based images

Other general-purpose OS
Other general-purpose operating systems

Misc utility images

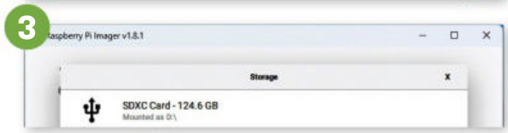

SDXC Card - 124.6 GB
Mounted as D:\

1. **Download and launch the Raspberry Pi Imager, selecting Raspberry Pi 5 as your device.**

2. **Next, pick 'Raspberry Pi OS (64-bit)' as your OS.**

3. **Then select your memory card as storage.**

4. **When asked if you'd like to apply OS customization settings, we recommend you click 'Edit settings'.**

5. **You can now give your Pi a network name, create a username and password, and enter Wi-Fi settings.**

6. **You might also want to enable SSH.**

7. **Finally, tell the Imager to write the OS to your card.**

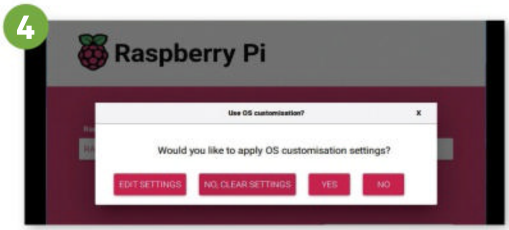

Would you like to apply OS customisation settings?

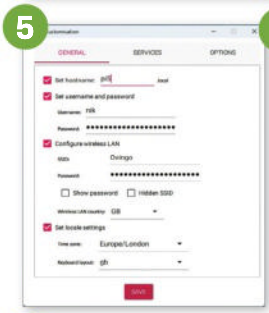

☑ Set hostname: pi5 .local
☑ Set username and password
Username: nik
Password: ********************
☑ Configure wireless LAN
SSID: Dvingo
Password: ********************
☐ Show password ☐ Hidden SSID
Wireless LAN country: GB
☑ Set locale settings
Time zone: Europe/London
Keyboard layout: gb

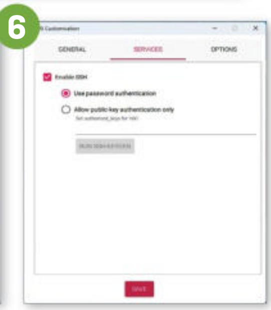

☑ Enable SSH
● Use password authentication
○ Allow public-key authentication only

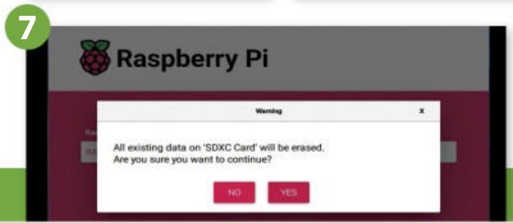

All existing data on 'SDXC Card' will be erased.
Are you sure you want to continue?

access to the GPIO pins. We also added the official Raspberry Pi snap-on active cooler for $5.18, to ensure desktop performance isn't affected by thermal throttling.

Like its predecessor, the Pi 5 has dual video outputs capable of driving a pair of 4K displays at 60Hz, so you can use more or less any monitor. Since it uses micro-HDMI ports, however, you might need a micro-HDMI to regular HDMI cable: you can pick up the official Raspberry Pi version for $5.18 through official resellers. Similarly, you can plug in any USB-compatible keyboard and mouse, or pay $29.72 for the admirably unfussy official Raspberry Pi desktop set.

There's one final component we'd love to add, but it's not yet available. The M.2 HAT will let you mount an NVMe SSD, bringing high-speed native storage to the Pi for the first time. This is expected to arrive by early 2024, and cost around $28.

All told, our own desktop setup, comprising an 8GB Raspberry Pi 5, Pibow Coupe 5 case, official power supply, active cooler, and micro-HDMI cable, came to $113.44. For the rest, we used an existing monitor, keyboard, and mouse, and recycled a 128GB microSD card as our Pi system drive.

Setting up Raspberry Pi OS

Setting up the Raspberry Pi 5 for desktop use isn't particularly difficult, although as usual you'll need an existing computer with a microSD card reader to create the boot media. Start by downloading the Raspberry Pi Imager from raspberrypi.com/software onto your PC, Mac, or Linux system, then insert your microSD card, launch the Imager tool, and click your way through the imaging process. To install the standard Raspberry Pi OS, select Raspberry Pi 5 as your device, pick 'Raspberry Pi OS (64-bit)' as your operating system, and your memory card as storage.

When asked if you'd like to apply OS customization settings, we recommend you click 'Edit settings'; you can now give your Pi a network name, create a username and password, and enter your Wi-Fi settings. Doing these things means you won't need to set them up when booting up the Pi. You might also want to click on the Services tab and enable SSH, as this gives you the option of opening a remote terminal on the Pi (even when it doesn't have a monitor connected), which can be handy for remote maintenance.

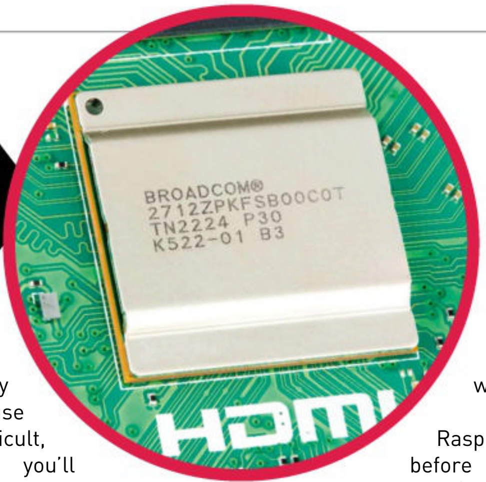

The Pi 5 uses a 2.4GHz Cortex-A76 CPU—a huge step up from the Pi 4.

Finally, tell the Imager to write the operating system to your card. The time this takes will be determined by the speed of your PC and the card: in our case, it took around five minutes with a Class 1 card.

If you've used Raspberry Pi OS before (or its predecessor Raspbian), the latest release should be familiar. However, one significant change is that the new OS uses the Wayland windowing system in place of X11, as it's more secure and efficient. You probably won't notice a difference in everyday use, but the change means that the RealVNC remote desktop server is no longer supported: the new Pi OS uses wayvnc, which you can remotely access with a compatible client like TigerVNC (tigervnc.org).

Alternatively, if you want to keep using RealVNC you can switch back to X11. To do this, open a Terminal window on the Pi 5 itself, or connect from Windows via SSH, and enter:

```
sudo raspi-config
```

Navigate to Advanced Options and press Return; then navigate to Wayland and press Return. Press Return again with option W1 selected, press OK, then select Finish, and allow the Pi to reboot.

Running Ubuntu on the Pi

The Raspberry Pi OS is tailor-made for the lightweight board, but if you prefer to use an industry-standard Linux desktop, that's no problem. The Pi can run a range of ARM-compatible distributions, including the popular Ubuntu platform.

You can set this up using the standard Raspberry Pi Imager tool. When choosing your operating system, click through to 'Other general-purpose OS', then select Ubuntu. You can choose between Desktop, Server, and Core (IoT) versions, and pick either the latest release or a more stable build with long-term support. The Imager tool can't set up your user account or networking configuration in the same way as it can with Raspberry Pi OS; you'll be dropped into the Ubuntu setup wizard when you boot the Pi from your newly imaged card.

©RASPBERRY PI FOUNDATION

WE DON'T SERVE YOUR TYPE

The free LibreOffice suite interoperates very cleanly with Microsoft Office, but if you're sharing files with Office users, you may hit formatting issues, because the Raspberry Pi OS doesn't include Microsoft's proprietary fonts. You can improve matters by installing the Microsoft Core TrueType Fonts pack, which contains 11 fonts including Arial, Verdana, and Times New Roman. To do this, open a Terminal window on your Pi (the quickest way is to press Ctrl+Alt+T) and enter the following two lines:

```
sudo apt update
sudo apt install ttf-mscorefonts-installer
```

Unfortunately, the bundle doesn't include Calibri or Cambria, which are the default fonts in Office, and Microsoft doesn't offer a legal way to install them on your Pi.

One workaround is to download Carlito and Caladea from Google Fonts instead. These typefaces are metrically identical to the Microsoft fonts, as well as being similar in appearance, so you can substitute them without affecting the formatting or look of a document. When sharing with Office users, you can embed Google fonts in the document, advise your colleagues on how to substitute them, or distribute your work as a non-editable PDF file.

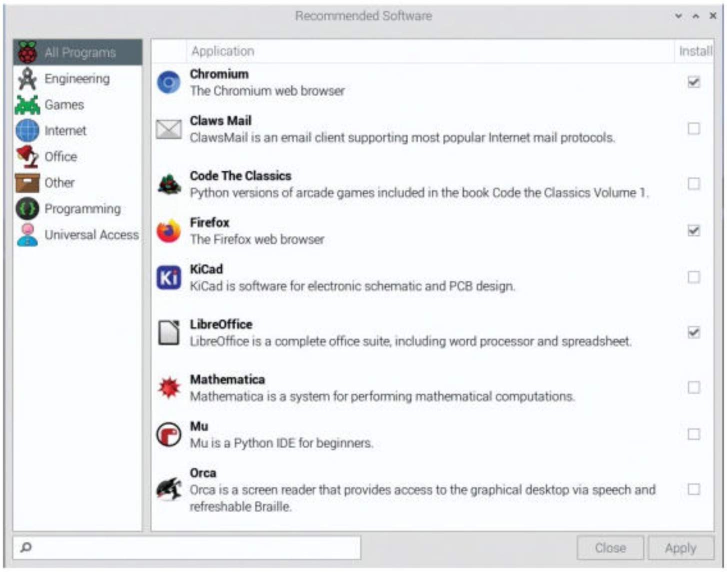

Chromium — The Chromium web browser ☑

Claws Mail — ClawsMail is an email client supporting most popular Internet mail protocols. ☐

Code The Classics — Python versions of arcade games included in the book Code the Classics Volume 1. ☐

Firefox — The Firefox web browser ☑

KiCad — KiCad is software for electronic schematic and PCB design. ☐

LibreOffice — LibreOffice is a complete office suite, including word processor and spreadsheet. ☑

Mathematica — Mathematica is a system for performing mathematical computations. ☐

Mu — Mu is a Python IDE for beginners. ☐

Orca — Orca is a screen reader that provides access to the graphical desktop via speech and refreshable Braille. ☐

Installing your applications

Once you've got your Pi set up and connected , the last step is to set up your applications. Raspberry Pi OS preinstalls only a minimal software set, so—unless you want to rely on web apps—you'll need to install some programs. Start by browsing Raspberry Pi's recommendations, found by opening the Raspberry menu, then selecting Preferences | Recommended Software.

We recommend LibreOffice for general productivity and GIMP for image manipulation; Inkscape is a powerful free vector graphics suite, and if you want to use your Pi for programming, Thonny is an excellent choice for working with Python and MicroPython. Each of these applications is also available for Windows and macOS, so your Raspberry Pi 5 can become part of a wider cross-platform workflow.

Raspberry Pi can run a wide range of software that has direct equivalents for both Windows and macOS.

Other programs can be installed from the Raspberry Pi software store, which you can access by clicking Add/Remove Software on the Preferences menu. For anything web-based, Raspberry Pi OS includes both Firefox and Chromium browsers by default, and you can naturally install others if you have accounts set up elsewhere. Needless to say, if you hit any problems or need advice, you can always turn to the huge, highly active Raspberry Pi and Linux communities for pointers. ⏻

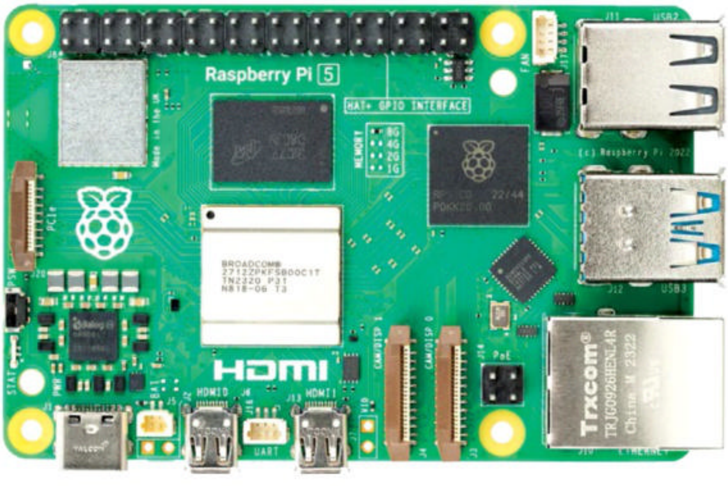

The Pi 5 provides some new interfaces, as well as a power button on the board.

HOW FAST IS FAST?

We've said that the Raspberry Pi 5 is fast enough to use as an everyday desktop machine, but what does that really mean? We used the Geekbench 6 benchmark to compare its performance to other low-cost computers.

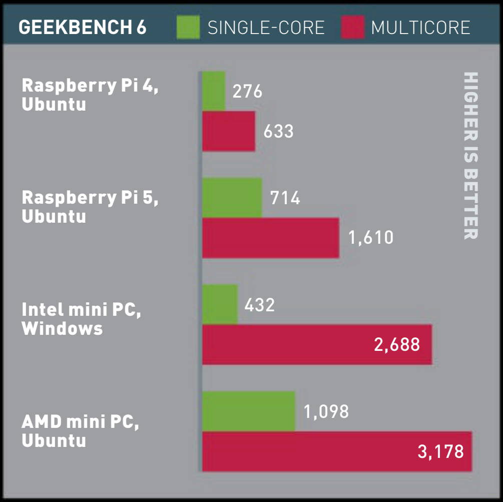

We started with a mini PC with 16GB of RAM, running Windows 11 Pro on an Intel N100 processor and costing around $215. We then repeated the test on a second mini PC, this time on a much meatier quad-core AMD Ryzen 5 system with 8GB of RAM, costing around $600 and running Ubuntu Linux.

Next, we tested the Raspberry Pi 4 Model B revision 1.4, with its 1.8GHz ARM processor and 8GB of memory—previously the most powerful model available, before the arrival of the Pi 5. We fitted an anodized aluminum heatsink on both sides to avoid processor throttling during our tests. Since there's currently no official ARM build of Geekbench for Linux, we used the preview build from geekbench.com/preview, running on the latest release of Ubuntu as recommended by the publisher. Finally, we ran the same test on a Raspberry Pi 5.

As you can see from the graphs above, the Raspberry Pi 5 is a huge step up from the Pi 4. For both single-core and multicore performance it's more than two-and-a-half times as fast as its predecessor, and it even beat the Intel PC for single-core operations. It was no surprise to see that the Raspberry Pi 5 couldn't keep up with the AMD system, but considering it costs less than a quarter of the price, it acquitted itself very creditably.

Perhaps more importantly, the Pi 5 feels fast in use. While using it we noted that web pages loaded snappily, and LibreOffice Writer launched in under four seconds. That's not far behind our Linux-based mini-PC, which took two seconds—and well ahead of the Windows test machine, which took 13 seconds to get LibreOffice Writer up and running.

CENTERFOLD

PERFORMANCE GEAR LAID BARE

Gigabyte Aorus Gen5 12000 SSD

TO BE AT the cutting edge requires sacrifice. It doesn't matter if that's in sport, life, or PC hardware. To operate at the very edge of what's possible is a give and take matter. It's not smart, clever, or the sensible thing to do, but it's not meant to be. To be a pioneer, early adopter, or explorer requires you to take risks.

Nowhere is there a better example of this than with Gigabyte's Aorus Gen5 12000 SSD, the company's second generation of its prestigious PCIe 5.0 M.2 SSDs. This glorious 2TB solid state drive packs in some serious firepower, with 3D TLC NAND Flash ripping and roaring along at incredible speeds, peaking at 12,400 MB/s on the read, and 11,800 MB/s on the write. It's blisteringly quick—22.54x faster than a traditional 2.5-inch SSD from yesteryear.

There is, as always, a caveat to that speed and power, and it's heat. Although the NAND can handle it, and the controller is primed and ready, the cooling solutions have needed a significant upgrade to handle the phenomenal work-rate these drives can tolerate. Gigabyte has paired this with its latest M.2 Thermal Guard XTREME heatsink, complete with twin heatpipes, plus twin thermal pads, and an aluminum M.2 baseplate to keep it cool, even under the most intense of systems loads.

Right now, this is one of—if not the—fastest SSD money can buy, without falling to gimmickry or RAID instability, and Gigabyte knows it. –ZAK STOREY

EAM UP. FIGHT ON
AORUS

1 THE ULTIMATE NAND FLASH

The Gen5 12000 is utilizing 232-layer 3D TLC NAND Flash. It's also equipped with 4GB of LPDDR4 external DDR cache, all operated by Phison's latest PS5026-E26 controller.

2 INSANE COOLING SOLUTIONS

Most PCIe 5.0 drives require a fan to keep their temperatures cool. The Gen5 12000 ignores that, and goes for a massive twin heatpipe heatsink, complete with Nanocarbon coating.

3 CHUNKY WARRANTIES

The 12000 is rated at 1.6 million hours before failure, or 1,400 TBW. On top of that, it features a five-year warranty, or whichever metric comes first.

Run Windows on your Apple Silicon Mac for free

Boot Camp may have got the boot, but it's still possible to run Windows applications on macOS. *Darien Graham-Smith* finds out how

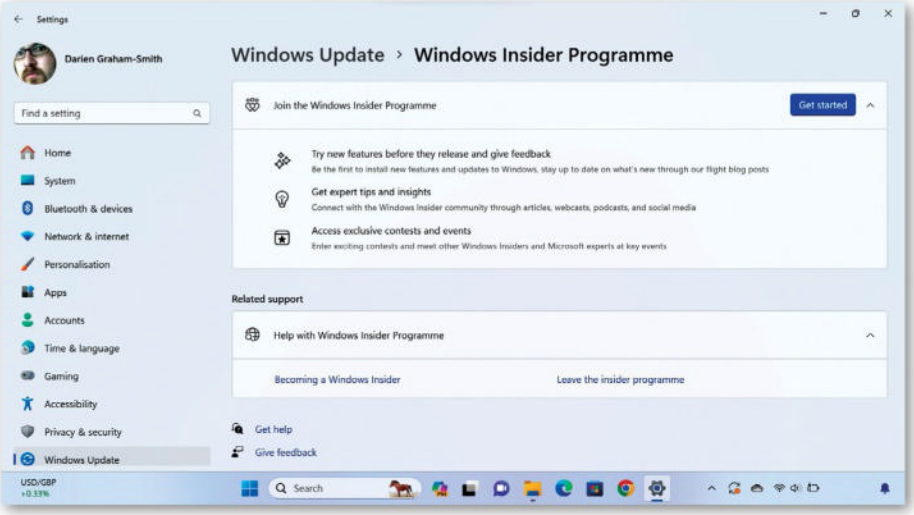

Windows Update › Windows Insider Programme

Join the Windows Insider Programme

Try new features before they release and give feedback
Be the first to install new features and updates to Windows, stay up to date on what's new through our flight blog posts

Get expert tips and insights
Connect with the Windows Insider community through articles, webcasts, podcasts, and social media

Access exclusive contests and events
Enter exciting contests and meet other Windows Insiders and Microsoft experts at key events

Related support

Help with Windows Insider Programme

You can join the Insider program from within Windows 11.

When Apple migrated the Mac platform from Intel CPUs to its own ARM-based chips in 2020, it spelt the end for the Boot Camp utility. Windows simply couldn't boot on the new M1 hardware, so the popular dual-booting tool was quietly ditched.

However, if you still want to run Windows applications on your Mac, there is a way. For more than a decade, Microsoft has been working on bringing Windows to the ARM platform, and its latest releases are now more than mature enough for day-to-day use. If you want to try out some cross-platform action on your own Apple Silicon Mac, you can buy an ARM edition of Windows 11 from the Microsoft online store (tinyurl.com/5c7t89rp)—or download the latest Insider Preview release and run it for free. Here's how.

Expectation management

The great thing about Windows on ARM is that it has a built-in translation layer, which enables it to run existing Intel applications without modification. The ARM version of Windows 10 only works with 32-bit code, but the latest Windows 11 builds support 64-bit software, too, so you can run almost any application on ARM hardware. The main exceptions are programs that try to interface directly with the hardware or install their own system drivers, such as 3D games or third-party security tools, although the built-in Windows Defender is still there to protect you against online attacks.

This translation layer does, however, incur a certain performance penalty. What's more, since Apple Silicon doesn't support dual booting, Windows has to run inside a virtual machine, meaning it can't make full use of your Mac's memory and CPU resources. The upshot is that Windows applications should still be usable, but may feel sluggish and jittery compared to what you're used to on your Mac.

Finally, bear in mind that Windows Insider builds come with no formal guarantee of stability or support. While you're not legally obliged to activate your Insider OS (see tinyurl.com/2p9hdpv3 for the full terms of the programme), a few desktop customization features won't be available unless you enter a valid product key.

In all, if you want to run Windows as your primary operating system, you'll probably get a better experience from a real Intel PC. But if you're happy using your Mac for everyday tasks, and just occasionally need to run a few specific programs in Windows, a virtual machine running Insider code could well be all you need.

Getting the Windows installer files

Right now, Microsoft doesn't publicly offer a downloadable installer for the 'stable' edition of Windows on ARM, unless you buy it through the online store. However, it's easy to get hold of, as we'll see below. Just remember that it's not legal to use the commercial edition without a licence and product key.

If you prefer to take the free route, you can download the Insider Preview installer from Microsoft, as long as you're a member of the Windows Insider Preview program. If you haven't joined, register at tinyurl.com/4ydpvnms. You can also join via the Settings app on any PC running Windows 10 or 11, but this will configure that PC to receive pre-release OS updates.

Once you've enrolled in the Insider Programme, you can download the Windows 11 installation files for ARM from tinyurl.com/2fwxjw63 (you may need to sign in to your Microsoft account again to confirm your Insider membership). The page will present a drop-down menu offering you three different builds of Windows 11: the Canary and Dev releases give you a chance to try the very latest code, while the Beta version should be the most stable. After making your choice, hit Confirm.

Next, you'll be prompted to choose your language—at present the only option is US English. After this, a download link should appear. The Windows 11 installer comes down the line in the form of a large (around 10GB) disc image file in Microsoft's VHDX format.

Having saved this file onto your Mac, you can cancel your membership of the Windows Insider Program. Visit tinyurl.com/fm2rtswk and click 'Leave the program now', or go back to the Windows Insider settings on your PC, scroll to the bottom, and click 'Leave the Insider programme'.

Choosing a host and preparing your installation media

MacOS doesn't come with a built-in virtualization host, but there

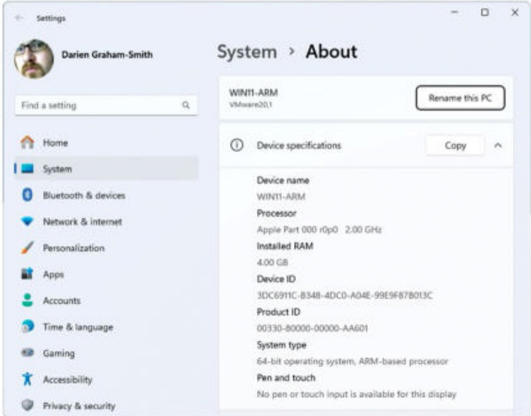

System › About

WIN11-ARM
VMware20,1

Device specifications

Device name
WIN11-ARM
Processor
Apple Part 000 r0p0 2.00 GHz
Installed RAM
4.00 GB
Device ID
3DC6911C-B348-4DC0-A04E-99E9F87B013C
Product ID
00330-80000-00000-AA601
System type
64-bit operating system, ARM-based processor
Pen and touch
No pen or touch input is available for this display

Windows happily recognizes and supports Apple Silicon hardware.

©MICROSOFT

How long can you use Windows 11 Insider Preview for?

When you first boot up Windows 11 on your Mac, you may be greeted by a banner telling you that your installation has expired. It seems Microsoft isn't updating its installation images that regularly, so the OS you've installed could be out of date on day one.

Don't panic—you can use Windows for as long as you like. You simply need to download the latest build via Windows Update, just as on an Intel-based PC. Even if you've left the Windows Insider Preview programme, your Insider OS will receive the latest patches and builds. However, there's currently no off-ramp; to leave the testing track and switch to release code, you'll need to carry out a clean reinstallation of the OS.

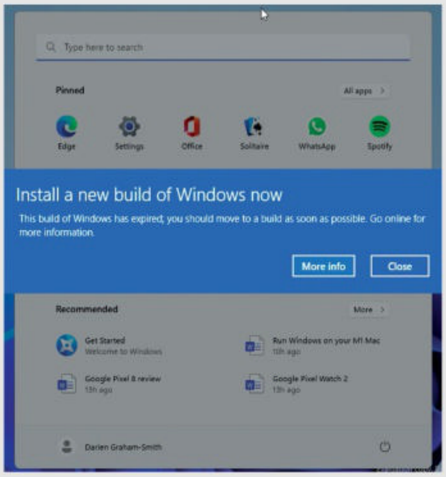

Install a new build of Windows now

This build of Windows has expired; you should move to a build as soon as possible. Go online for more information.

are plenty of third-party tools to choose from. The slickest is Parallels (parallels.com), which supports some neat features, such as running Windows applications directly on the Mac desktop. However, it costs $130 including VAT for personal use—and that doesn't include a Windows licence. The best way to buy it is as a bundle, with Parallels often being bundled with the ARM version of Windows 11 Pro.

If you're looking for a free solution, the open-source UTM hypervisor is a popular choice (mac.getutm.app), but it has limited support for saving and resuming VM snapshots, so it's not ideal for quickly firing up Windows. Another option is Oracle VirtualBox (tinyurl.com/yc583tdf); this hasn't yet been officially released for Apple Silicon, but you can try out a Developer Preview of version 7.0.6 from the link above.

Our preferred free choice is VMware Fusion Player 13.5, a full, mature virtualization host that's free for personal use: download it from tinyurl.com/yun5usay. The catch is that VMware doesn't natively support VHDX files, so if you've downloaded your Windows 11 installation media in that format, you'll need to convert it into VMware's own VMDK format. But that's not hard: open a Terminal window on your Mac and, assuming your VHDX file is in your Downloads folder, enter these two commands:

```
cd ~/Downloads
/Applications/VMware\ Fusion.app/Contents/Library/vmware-vdiskmanager -r Windows11_InsiderPreview_Client_ARM64_en-us_22598.VHDX -t 0 Windows11.VMDK
```

If your downloaded VHDX file has a different filename to ours then adjust the command appropriately. In a few moments you should have a new VMDK version of the installer file sitting in your Downloads folder.

Creating the VM

When you first open VMware Fusion Player, you'll see a window headed 'Select the Installation Method'. If you want to download the commercial release of Windows 11, click 'Get Windows from Microsoft' and hit Continue. On the next page, leave 'Windows Edition' set to 'Professional' (this also supports the Home edition), select 'English (en-us)' as your language, and hit 'Download Windows'.

If you're installing from the Insider Preview image, drag the VMDK file you just created into the installation window. VMware Fusion Player will detect the platform and architecture of your installer (Windows 11 64-bit ARM), so hit Continue.

From here, the setup procedure is largely the same, whether you're using your own VMDK or an automatically downloaded image. The next page asks you to confirm UEFI BIOS emulation—do this by hitting Continue—then invites you to set up a password for the virtual TPM inside your virtual Windows machine. Click 'Auto Generate Password' if you wish; it will be saved in your macOS Keychain if you ever need to refer back to it (for example, if you want to move your VM to another Mac or access it from a different user profile).

Hit Continue to move onto the virtual disk setup page. If you're using the official installer downloaded by VMware, leave 'Create a new virtual disk' selected and hit Continue. The default size is 64GB, but the image will only take up as much disk space as is needed to hold the files used by Windows—and if you need more space you can always expand it later.

To install the Insider Preview, select 'Use an existing disk', and navigate to your VMDK file in the requester that opens. Leave 'Make a separate copy of the virtual disk' selected so that you still have a clean copy of the installer.

Hit Continue, and you'll see a summary of your new VM configuration. Click 'Customize Settings' to tweak the virtual hardware; the default settings will assign 4GB of RAM and two processor cores to your Windows VM—fine for lightweight desktop use, but you can increase these allocations. Remember that it's a balancing act, as these resources won't be available to macOS while the VM is running.

Hit Finish, and VMware will prompt you to save your new VM. You'll then see a black screen with a big 'Play' symbol on it. Click this—or the little icon at the upper left of the title bar—to 'power on' your virtual machine.

Installing Windows

Turning on the VM brings up a simulated BIOS screen. If you're

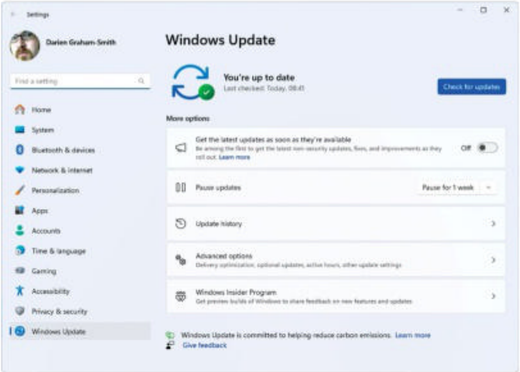

Windows Update

You're up to date
Last checked: Today, 08:41

Get the latest updates as soon as they're available
Be among the first to get the latest non-security updates, fixes, and improvements as they roll out. Learn more

Advanced options
Delivery optimization, optional updates, active hours, other update settings

Windows Insider Program
Get preview builds of Windows to share feedback on new features and updates

Windows Update is committed to helping reduce carbon emissions. Learn more

Keep your ARM edition of Windows updated in the usual way.

using the full release version of Windows 11 on ARM, you'll see a message saying 'Press any key to boot from CD or DVD…'—and you should do this, as the installation files are on the virtual DVD drive.

You'll then see the legacy Windows installer interface. This will prompt you with a few basic questions, such as your region and keyboard preferences, which edition of Windows you want to install, and which hard disk you want to install it on. Once you've started the installation process, the system will copy the files, reboot, spin for a few minutes, then launch the graphical setup wizard. If you're using the Insider Preview, the installation files are already in place, so your VM should boot directly into the setup wizard.

Once you see the opening screen of the wizard, you might want to start working through it. Try this, though, and you'll run into difficulties, as the installer needs network access, which in turn requires us to install a driver.

Open the Virtual Machine menu and select 'Install VMware Tools' (or 'Update VMware Tools'); on the requester, hit 'Install' to mount the disc containing the necessary drivers. Next, press Shift+F10 to bring up a Command Prompt—if that doesn't work, press Cmd+R to bring up the 'Run…' requester, and enter 'cmd'. Now enter:

```
d:setup.exe
```

This will launch the installer for VMware's tools and drivers. Click through it, accepting the default options, and reboot the virtual machine when prompted. When the installer comes back up, you'll be able to complete the installation without any technical hurdles.

The installer will ask you to name your device—something like WIN11-ARM should be nicely unambiguous—and whether you want to set it up for personal use, work, or school. Select the former, as this virtual machine isn't on a corporate network, and log in with your Microsoft account.

The installer will also offer to bring across your OneDrive files, preferences, and apps from your most recently used Windows device; if you prefer a clean installation, click 'View more options' and choose 'Set up as a new device'. You'll be asked to permit all sorts of personalization services; it's a bit tiresome, but you only have to go through it once.

The installer will spend a few more minutes building your user profile, and you'll be greeted by the familiar Windows desktop. What you do from here is up to you: you may well want to start by installing your commonly used programs.

You can also tweak your VMware Fusion Player settings. Dragging the edges of the window will dynamically resize your Windows desktop; display scaling can either be adjusted natively in Windows, or in VMware's System Settings | Display page. At present there's no way to set up native folder sharing between your Windows and Mac, but you can work around this by sharing folders via the network. The VM can be shut down, suspended, or rebooted via the Virtual Machine menu, and you can revert to a previous state via Snapshots.

For complete documentation on using VMware Fusion Player to run Windows 11 on Apple Silicon hardware, check out tinyurl.com/ywwbnzex—and you may also want to look over the unofficial guide at tinyurl.com/4vk9k9cb.

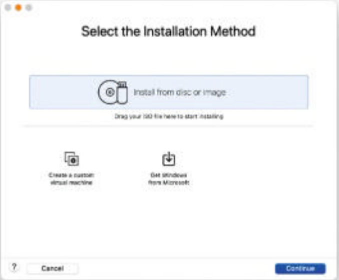

Select the Installation Method

Drag your ISO file here to start installing

Install from disc or image
Create a custom virtual machine
Get Windows from Microsoft
?
Cancel
Continue

Open your converted image file, or download installation media from Microsoft.

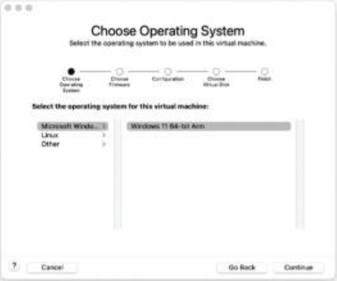

Choose Operating System

Select the operating system to be used in this virtual machine.

Select the operating system for this virtual machine:

VMware recognizes your OS, so you can click straight through.

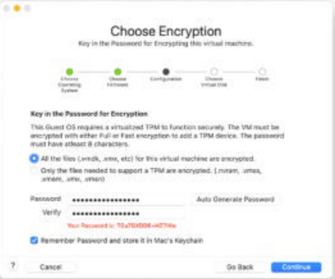

Choose Encryption

Key in the Password for Encrypting this virtual machine.

Key in the Password for Encryption

This Guest OS requires a virtualized TPM to function securely. The VM must be encrypted with either Full or Fast encryption to add a TPM device. The password must have atleast 8 characters.

● All the files (.vmdk, .vmx, etc) for this virtual machine are encrypted.

○ Only the files needed to support a TPM are encrypted. (.nvram, .vmss, .vmem, .vmx, .vmsn)

Password ••••••••••••••••

Verify ••••••••••••••••

☑ Remember Password and store it in Mac's Keychain

You'll need to set a password for TPM emulation.

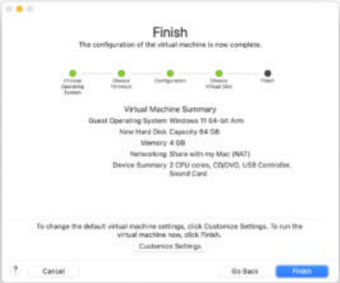

Finish

The configuration of the virtual machine is now complete.

Virtual Machine Summary

Guest Operating System Windows 11 64-bit Arm

New Hard Disk Capacity 64 GB

Memory 4 GB

Networking Share with my Mac (NAT)

Device Summary 2 CPU cores, CD/DVD, USB Controller, Sound Card

To change the default virtual machine settings, click Customize Settings. To run the virtual machine now, click Finish.

Tweak the virtual hardware settings to suit your needs, or go with the defaults.

What next for Windows on ARM?

The first ARM-based version of Windows was Windows RT, created in 2012 for the original Microsoft Surface tablet. It looked almost identical to Windows 8, but it only supported ARM-native code, and couldn't run existing Windows desktop apps. Unsurprisingly, it flopped.

Thankfully, Microsoft learnt from that experience, adding an application translation framework to the ARM editions of Windows 10 and 11. Now that ARM PCs can run regular Windows software, the platform is working its way into the mainstream: Microsoft's Surface Pro 9 convertible is built on ARM silicon, as are a number of third-party Windows laptops.

Even so, ARM remains a second-class platform in the Windows world. The chips in those systems can't match the speeds of the latest processors from AMD and Intel, and running applications under translation doesn't help with performance or power.

That could change, though. Qualcomm just announced its new generation of 'Snapdragon X' ARM processors, promising to bring Mac-rivalling performance and battery life to the Windows platform in 2024. There's more innovation and competition on the way, too, with both AMD and Nvidia reportedly working on rival ARM chips for 2025.

We can expect to see a lot more ARM-based PCs and laptops in the next few years. As the market grows, developers will be incentivized to offer high-performance ARM-native versions. After more than a decade in the doldrums, Windows on ARM may get its chance to shine. You can look forward to the benefits, even on your Mac.

© MICROSOFT

gamesradar+

One-stop destination for games, movies, TV, and more!

New Channels: Movies, TV and Cool Stuff!

www.gamesradar.com

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

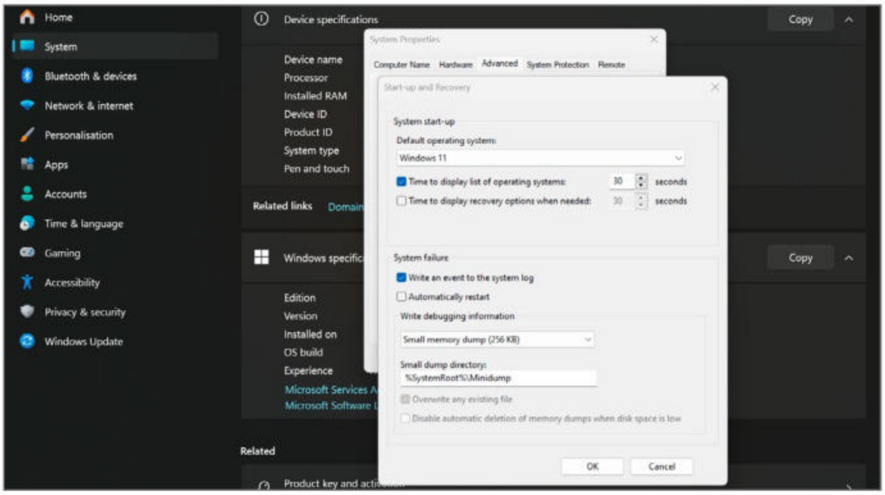
Home
System
Bluetooth & devices
Network & internet
Personalisation
Apps
Accounts
Time & language
Gaming
Accessibility
Privacy & security
Windows Update
Device specifications
Copy
Device name
Processor
Installed RAM
Device ID
Product ID
System type
Pen and touch
Related links
Domain
Windows specific
Copy
Edition
Version
Installed on
OS build
Experience
Microsoft Services A
Microsoft Software L
Related
Product key and activation
System Properties
Computer Name
Hardware
Advanced
System Protection
Remote
Start-up and Recovery
OK
Cancel

System start-up

Default operating system: Windows 11

☑ Time to display list of operating systems: 30 seconds

☐ Time to display recovery options when needed: 30 seconds

System failure

☑ Write an event to the system log

☐ Automatically restart

Write debugging information: Small memory dump (256 KB)

Small dump directory: %SystemRoot%\Minidump

☑ Overwrite any existing file

☐ Disable automatic deletion of memory dumps when disk space is low

STOP WINDOWS RESTARTING ON CRASH

Ever shut down your PC and gone to bed, only to find your rig back on in the morning? Want a fix? Type in 'settings' in the Start menu, then go to 'System', then 'About'. Click 'Advanced System Settings', make sure you're in the 'Advanced' tab, then click 'Settings...' in 'Start-up and Recovery', and make sure 'Automatically restart' is unchecked under 'System failure'. Hit ok, click out, and you'll be good.

MAKE – USE – CREATE

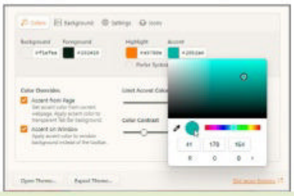

60
Browse the web your way with Vivaldi

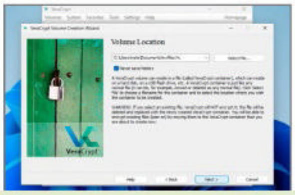

64
Secure your secrets with VeraCrypt

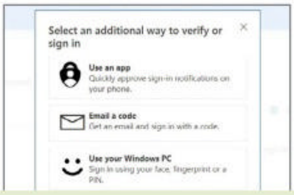

68
Use Bitwarden to store your Passkeys

© MICROSOFT, VIVALDI, VERACRYPT, BITWARDEN

ZAK STOREY
CONTRIBUTOR

GAMING CHAIRS, WHAT? WHY?

I'm in the market for a new office chair. Mine is dying, and it's not a pretty sight. The faux leather material has chipped off at the base, the foam stuffing has fallen apart, and the internal strapping has all but given up its hold on the seat base—it's just looking a little rough around the edges. It's a shame, because I've only had this since the start of the pandemic, and I expected it to last longer.

Admittedly, I put these things through a bit more wear and tear than the average Joe. I've worked from home pretty much continuously from early 2020 to late 2023, and I spend most of my time sat cross-legged on it (it's more comfortable for me).

Which brand do I go for? I could choose something ergonomically brilliant, such as noblechairs' Hero line (they have a *Warhammer 40K* one, which I love), but I just don't find them comfortable. That's always been noble's unofficial tagline; 'Ergonomic doesn't mean comfortable'. But here's the thing: you can't convince me otherwise. I don't care if it comes with the Nobel Prize for ergonomic design; if it's not comfortable to sit in for long periods of time, it's not a good chair. We only live once, I could get hit by a falling RTX 4090 tomorrow and snuff it. What good will my ergonomic chair be then?

↘ submit your How To project idea to: editor@maximumpc.com

Browse the web your way with Vivaldi

YOU'LL NEED THIS

VIVALDI
https://vivaldi.com
web browser

GIVEN HOW MUCH TIME we spend on the web, whether it's shopping, research, or just following our favorite sports team as they head towards another crushing defeat, it's criminal that we don't pay much attention to our web browser. If you're looking to maximize your web browsing experience, you need to investigate Vivaldi. With new features being added all the time, its claims of being the world's most customizable browser are well-founded (helped by ready access to the Chrome web store, allowing you to add all manner of extensions and other plugins).

In this tutorial, we'll investigate how Vivaldi can be set up to deliver on its promise. We'll focus on its main tab management tools that bring order to chaos and allow you to manage dozens of tabs without losing your place, and help you tap into other features, like its mail client and quick commands feature, to speed up using the app. We'll also reveal how you can set up the user interface how you like it, with all of Vivaldi's controls and tools positioned where you find them easiest to access, all wrapped up in an attractive, but readable skin. **–NICK PEERS**

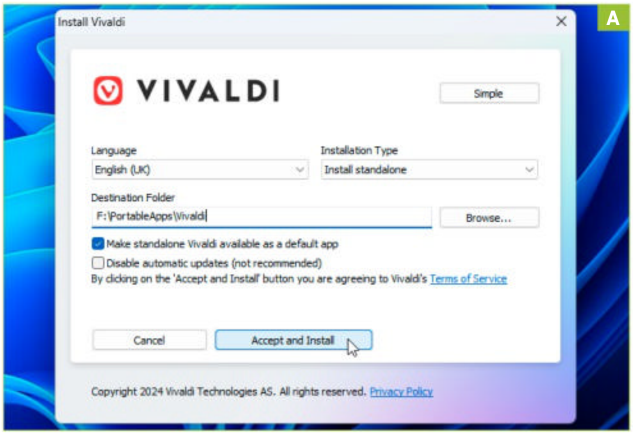

Language: English (UK)
Installation Type: Install standalone
Destination Folder: F:\PortableApps\Vivaldi

☑ Make standalone Vivaldi available as a default app
☐ Disable automatic updates (not recommended)

By clicking on the 'Accept and Install' button you are agreeing to Vivaldi's Terms of Service

Copyright 2024 Vivaldi Technologies AS. All rights reserved. Privacy Policy

1

FIRST TWEAK

If you're installing Vivaldi for the first time, your first opportunity to customize arrives the moment you run the installer. By default, Vivaldi installs into your user folder under AppData\Local—this is normal for a single-user install, but if you plan to install it for multiple users, click the Advanced button to reveal more options. You'll see an 'Installation Type' drop-down menu—select 'Install for all users' here to install to the usual Program Files folder. You'll also see an 'Install standalone' option. Select this, then click Browse to install Vivaldi to your portable apps folder. You can still make Vivaldi your default browser app (just leave the appropriate box checked as shown in [**Image A**]), but we suggest you leave 'Disable automatic updates' unchecked as recommended to ensure the portable app can update itself.

2

PICK A THEME

Vivaldi's appearance is uniquely customizable, but if you're in a hurry, then simply switch to a different theme—four are offered during the first-run wizard, but you can unlock many more. Open the Vivaldi menu and select Settings, followed by Themes. You'll see 10 options, as well as a 'New Theme...' option for creating your own, along with a 'Get more themes' link to visit the https://themes.vivaldi.net page, where you'll find hundreds more.

» Use the navigation tools to find a theme you like, then click Download. The theme will be downloaded and added to your collection, then Vivaldi will preview the theme by switching to it temporarily—click Install to keep it (and switch to that theme), or Cancel to remove it.

» The Editor button next to Library allows you to customize an existing theme (select it from the dropdown at the top of the page, as shown in [**Image B**]), or create a new one based on the currently selected theme (click New Theme under Library). Once created, rename the theme, then use the tools on the four tabs to edit it.

» Colors is where you set foreground, background, highlight, and accent colors, plus choose whether to follow the system theme (check 'Prefer System Color') and override accent colors with the current webpage's color. Background lets you choose your background image (which predominantly appears on Vivaldi's Start page), while Settings lets you set transparency, blur, and rounded corner settings. Finally, Icons makes it possible to replace the default Vivaldi icon set with another. You can drag your own choice of icon (28x28 pixels) on to each individual icon in turn, or plunder the icons from an existing theme. To do this, create a copy of the theme in question, then modify it to suit.

» After creating or choosing a theme, scroll down the Settings page beyond the Themes section to reveal more

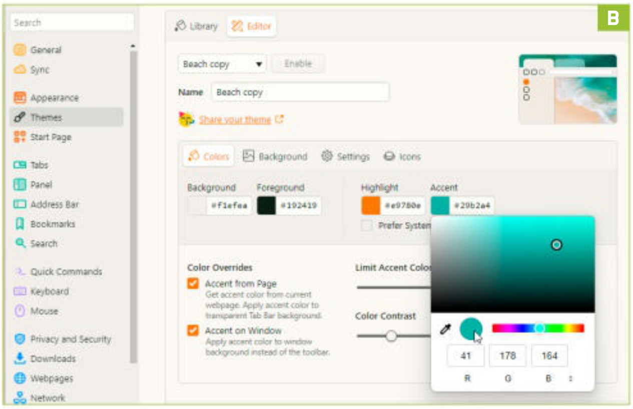

Beach copy
Name: Beach copy

Background: #f1efea
Foreground: #192419
Highlight: #e9780e
Accent: #29b2a4
☐ Prefer System

Color Overrides
☑ Accent from Page
Get accent color from current webpage. Apply accent color to transparent Tab Bar background.
☑ Accent on Window
Apply accent color to window background instead of the toolbar.

41 178 164

© VIVALDI

CUSTOMIZE THE START PAGE

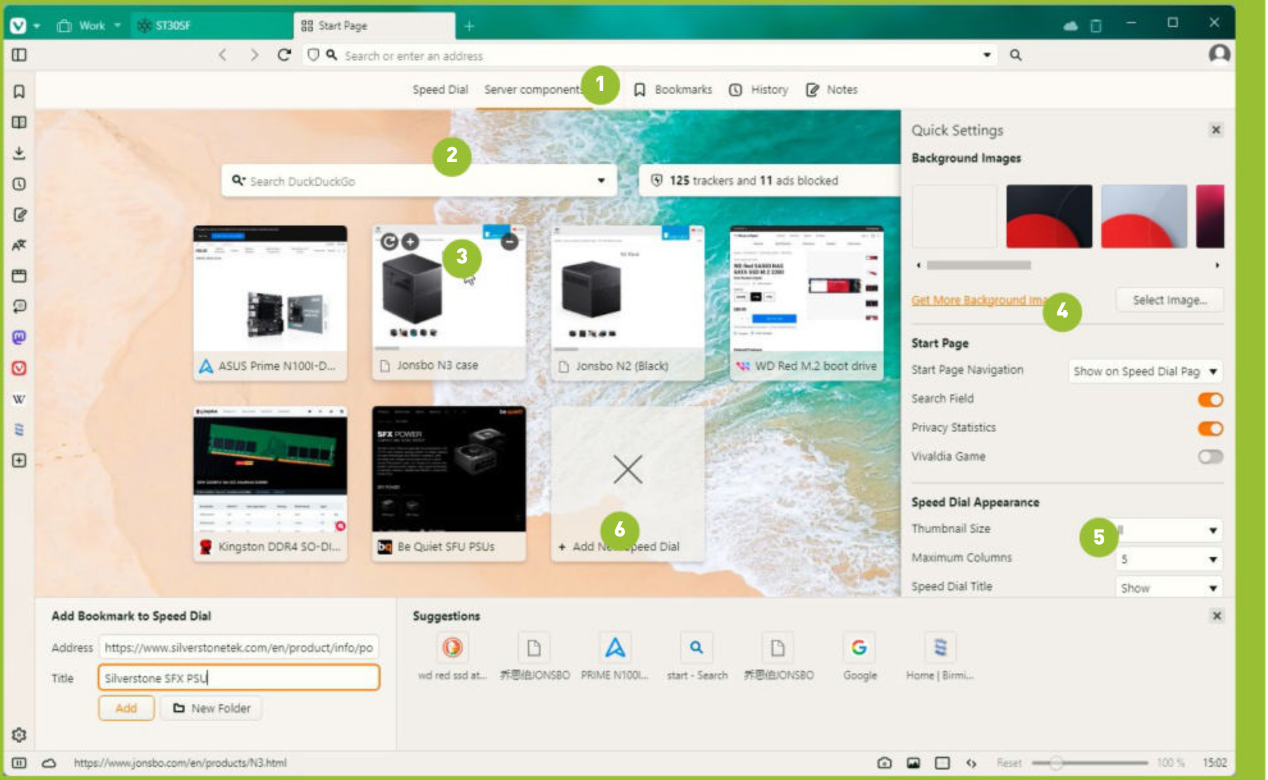
Work
ST30SF
Start Page
Search or enter an address
Speed Dial
Server components
Bookmarks
History
Notes
Search DuckDuckGo
125 trackers and 11 ads blocked
ASUS Prime N100I-D...
Jonsbo N3 case
Jonsbo N2 (Black)
WD Red M.2 boot drive
Kingston DDR4 SO-DI...
Be Quiet SFU PSUs
+ Add New Speed Dial
Quick Settings
Background Images
Get More Background Images
Select Image...
Suggestions
wd red ssd at...
乔思伯JONSBO
PRIME N100I...
start - Search
乔思伯JONSBO
Google
Home | Birmi...
https://www.jonsbo.com/en/products/N3.html
Reset
100 %
15:02
1
2
3
4
5
6

Start Page

Start Page Navigation: Show on Speed Dial Pag

Search Field: on

Privacy Statistics: on

Vivaldia Game: off

Speed Dial Appearance

Thumbnail Size:

Maximum Columns: 5

Speed Dial Title: Show

Add Bookmark to Speed Dial

Address: https://www.silverstonetek.com/en/product/info/po

Title: Silverstone SFX PSU

Add | New Folder

1. ADD SPEED DIAL FOLDERS
Too many favorites to list on a single Speed Dial page? Click the + button to add a subfolder inside the Speed Dial. Move existing entries to the new folder by right-clicking one and choosing Edit.

2. OPTIONAL ELEMENTS
Both Search bar and privacy stats can be removed from the Start page if you don't want them—open Quick Settings to do so.

3. SPEED DIAL CUSTOMIZATION
Update or replace a Speed Dial's icon or thumbnail by rolling your mouse over it and clicking one of two buttons in the top left. Rearrange Speed Dials using drag and drop.

4. QUICK SETTINGS PANE
Click the cog button in the top right corner to show this panel. Change the Start page's background image, plus customize (or hide) the bar at the top.

5. MORE DIAL CUSTOMIZATION
Use the Speed Dial Appearance section to access universal customization options like changing the thumbnail size and setting a limit on the number of columns shown.

6. ADD NEW SPEED DIAL
Add a new Speed Dial by clicking the + button to reveal this popup. Simply type the web address, give it a suitably descriptive title, and click Add.

theme-related options, such as choosing a different theme for private windows, setting up a theme schedule to rotate between themes, and even integrating Razer Chroma devices and Philips Hue lighting into your Vivaldi setup.

3 CHANGE SCREEN LAYOUT

Another concept introduced during the first-run wizard is the ability to choose where the various elements that make up Vivaldi's user interface go. The wizard lets you choose whether to place your tabs at the top, bottom, or either side of the main window, but that's just the tip of the iceberg.

» Again, changes are made in the Settings pane in the 'green' section containing Tabs, Panel, and Address Bar. Choosing to place the tabs on the left or right of the screen takes up more screen real-estate, but you'll see that the Tabs section offers many other customization options too, all largely self-explanatory. You'll see you can control exactly where new and duplicated tabs appear in the tab list, for example, set minimum width for tabs as their number grows on-screen, and set various other display options, like what kind of thumbnail is shown or where to display the close tab button. Scroll down further, and you'll discover features like tab stacking and workspaces—more on those shortly.

» Similar display options are available for the web panels, although here you're restricted to displaying them on either the left (default) or right of the screen. Finally, the Address Bar section can be placed at the top or bottom of the screen, regardless of where you placed your tabs bar.

» One element that can't (yet) be repositioned is the status bar at the bottom of the screen. You can, however, customize what's shown on it by right-clicking and choosing 'Edit > Customize Toolbar' to bring up the Toolbar Editor, as shown in [Image C], an option you'll also find on the Panel and Address Bar, too.

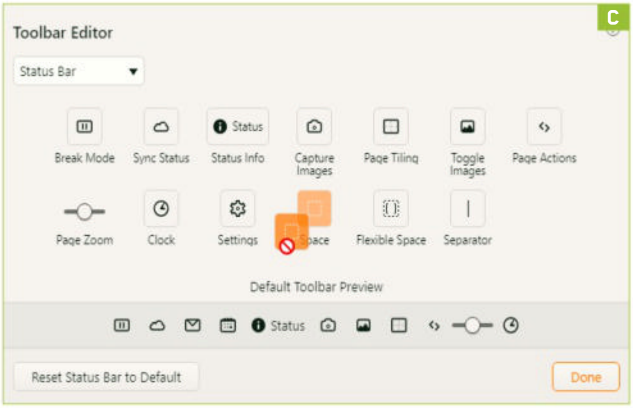
C
Toolbar Editor
Status Bar
Break Mode
Sync Status
Status
Status Info
Capture Images
Page Tiling
Toggle Images
Page Actions
Page Zoom
Clock
Settings
Space
Flexible Space
Separator
Default Toolbar Preview
Status
Reset Status Bar to Default
Done

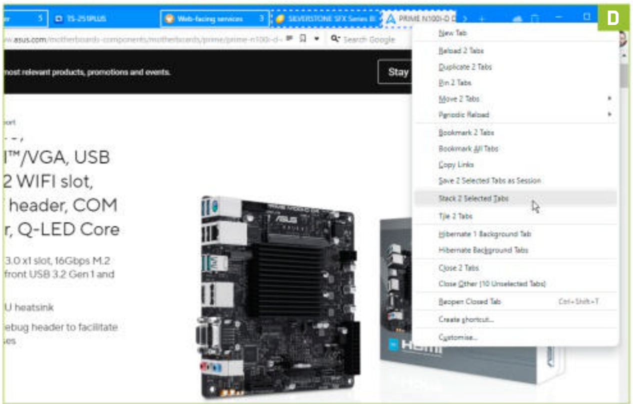
D
PRIME N100I-D
Stay
New Tab
Reload 2 Tabs
Duplicate 2 Tabs
Pin 2 Tabs
Move 2 Tabs
Periodic Reload
Bookmark 2 Tabs
Bookmark All Tabs
Copy Links
Save 2 Selected Tabs as Session
Stack 2 Selected Tabs
Tile 2 Tabs
Hibernate 1 Background Tab
Hibernate Background Tabs
Close 2 Tabs
Close Other (10 Unselected Tabs)
Reopen Closed Tab Ctrl+Shift+T
Create shortcut...
Customise...

4 STACK YOUR TABS

One area where Vivaldi stands head and shoulders above the competition is its tab management tools. Anyone juggling dozens of tabs will find that it offers numerous ways in which you can bring some order to the chaos and make it feasible to keep tabs open while hiding them from view.

» This is done using stacks and workspaces, which you can mix and match as required. Stacked tabs group the selected tabs into a single entry on the tab bar. Clicking this reveals a second row containing the stack's tabs. Simply hold Shift as you click on the tabs you wish to combine, then right-click and choose 'Stack 2 Selected Tabs', as shown in [**Image D**]. You can also instruct Vivaldi to automatically stack tabs pointing to the same base domain (such as pcgamer.com)—just right-click on any tab and choose 'Stack Tabs by Hosts' to do so.

» Once a stack has been created, it works the same as regular tabs—drag and drop tabs within a stack to rearrange them, plus drag tabs into and out the stack. Right-click the tab containing the stack, and you'll reveal more options, such as being able to reload all tabs within the stack (tabs are frozen when not directly accessed to reduce resource usage). You can also give the stack a different name and move the tab stack within the current set of tabs, to a new window, or even to another workspace.

5 WORKSPACES AND TILES

Another way to organize your tabs is into separate workspaces. Click Workspaces next to the main tab and it'll give you a succinct description of what it does, namely allow you to separate work from home life, or split your interests into distinct sections. Click New Workspace to give your new workspace a name, then click Create.

» You'll immediately switch to that workplace, where you'll see that it's separate from your original set of tabs (which sit above all workspaces until you move them). You'll also see that each workspace is allocated its own icon—click this from the Workspace dropdown menu to choose a different icon, or select an emoji instead.

» Once a workspace is created, you can move tabs and stacks to it by right-clicking them in their current workspace and choosing 'Move Tab Stack > Workspace > target workspace name'.

» Once two or more workspaces are set up, click the Workspaces dropdown menu to reveal options for rearranging them (via drag and drop), as well as renaming, deleting, and copying all the links to the clipboard as text URLs, as shown in [**Image E**]. You may see an option to hibernate the workspace, which also appears on select tabs and tab stacks. This occurs when there are 'active' pages consuming RAM and CPU resources; hibernating them simply puts them to sleep until you next access them.

» One final tip: if you frequently find yourself switching between two or more tabs, and you have a wide enough screen to accommodate them all, you can also 'tile' tabs so they're displayed next to each other in columns. To do this, simply hold Shift and click on each tab, then right-click and choose 'Tile x tabs' (where 'x' is the number of tabs selected).

6 SET UP PANELS

Another of Vivaldi's powerful features is its Panel, placed on the left of the Vivaldi window by default. The side panel itself isn't new—other browsers utilize it to lesser degrees—but it's packed full of useful tools you can switch on or off as required. You'll find panels covering bookmarks, reading lists, downloads,

MAIL AND CALENDAR

Unless you were looking for a replacement email client and calendar tool, chances are you skipped the opportunity to set them up during the initial wizard. Should you want to add these back in—plus access Vivaldi's RSS feed editor—navigate to Settings > Mail and check 'Enable Mail, Calendar and Feeds'. You'll see three new icons added to the Panel, where you can access each tool. You'll find that they're not token offerings, but powerful and flexible tools in their own right.

Before doing so, however, take the time to set up your mailboxes, link in a calendar account, and subscribe to some RSS feeds from the relevant parts of Settings. Once configured, the mail client works like any standalone application, and you'll see that the Mail section of Settings is populated with a wide range of options for formatting, composing, and being notified of new mail among other things—it even includes a contacts tool.

Similarly, the calendar tool syncs with most major providers or you can choose 'Local' to set one up from scratch that's stored locally only, or 'Vivaldi.net' to sync in with your Vivaldi profile. Again, once set up, the Calendar settings pane expands to show a wide range of tools and tweaks.

and history, but there's also a Notes panel for jotting down notes (and storing text from web pages). You'll also find a handy translation tool, and window manager, which thanks to its drag-and-drop capabilities makes it easy to move tabs between stacks, workspaces, and even windows. A recently added panel—Sessions—allows you to back up and restore different browsing sessions with ease (see the box).

» Beneath these come three examples of web panels, which basically reveal how you can pin any website as a panel for convenient access. There's a Mastodon panel for plugging in your Mastodon social media account, plus panels pointing to Vivaldi's Help pages and Wikipedia. Look closely, and you'll see that Vivaldi loads mobile-optimized version of the web pages to suit the panel's smartphone-like dimensions.

» Beneath this is a + button—click this, and you can add your own custom web panel. Type your target URL into the Address bar, and click + again to load the web page, which will load the mobile-optimized version by default and attempt to add the website's smart icon to the panel for easy identification. Right-click this to reveal more options, from opening the site in a new tab to loading the desktop version and making the panel float rather than remain pinned to the main Vivaldi window.

7 WORK WITH MULTIPLE SEARCH ENGINES

Like all good browsers, Vivaldi supports multiple search engines, as evidenced by the search bar. But why waste good real estate with a separate search bar just to be able to choose different engines when you can set up shortcuts for them within the main Address bar itself?

» To do this, navigate to Settings > Search, where you'll find the usual flurry of options giving you full control over how search works. Start by selecting 'Show as a Button' beneath 'Search Field Display' to reduce the box to a clickable button, so it takes up less room. Then scroll down to the Search Engines section, where you'll see all currently installed search engines, along with a handy one-letter nickname (such as 'b' for Bing, or 'd' for DuckDuckGo).

» Typing this letter followed by a space and your search terms opens the relevant website, but check 'In Address Field' under 'Allow Search Suggestions' along with 'Only When Using Search Engine Nickname', and you'll get suggestions from the search engine, too.

» You'll also see an option checked called 'Enable Direct Match'—uncheck this if privacy is of paramount importance, but ultimately it can get you to web pages quicker, earns Vivaldi a small amount of revenue to help support the project, and promises to come with no tracking information involved.

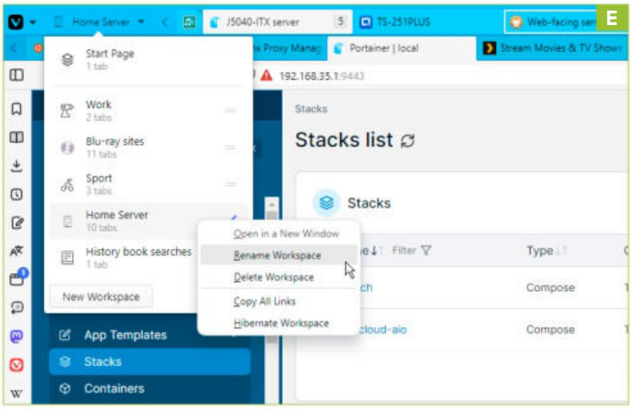
E
Home Server
J5040-ITX server
TS-251PLUS
Web-facing ser
Proxy Manag
Portainer | local
Stream Movies & TV Shows
192.168.35.1:9443
Start Page
1 tab
Work
2 tabs
Blu-ray sites
11 tabs
Sport
3 tabs
Home Server
10 tabs
History book searches
1 tab
New Workspace
Open in a New Window
Rename Workspace
Delete Workspace
Copy All Links
Hibernate Workspace
App Templates
Stacks
Containers
Stacks
Stacks list
Stacks
Filter
Type
Compose
Compose
cloud-aio

© VIVALDI

WORKING WITH SESSIONS

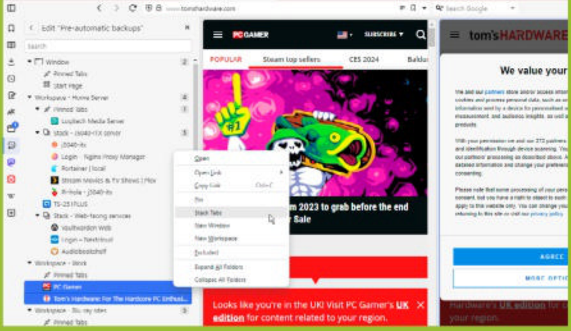
Edit "Pre-automatic backups"
Window
Workspace - Home Server
PC GAMER
tom'sHARDWARE
We value your
Looks like you're in the UK! Visit PC Gamer's UK edition for content related to your region.

Sessions serve two purposes in Vivaldi: their primary purpose is provide backups of your open tabs, stacks, and workspaces, but they can also be used to allow you to set up multiple configurations for different purposes, allowing you to switch between them with ease.

A recent addition to Vivaldi is the Sessions manager, accessible via the side pane. This makes it far easier to create and edit these sessions, making them painless to manage. From here, you can click + to back up the current session—give it a name, choose whether to include all workspaces or windows, and click Save.

The session will appear in the Sessions pane—to tweak it, select it from the list, and click Edit to the right of the +/- buttons. The contents of the saved session will appear in a hierarchical tree in the same way the current set appears in the Windows pane. The Sessions pane works in a similar way, allowing you to rearrange tabs by drag and drop, plus pin, stack, and unstack tabs via the right-click menu.

There are some limitations, likely to be lifted in a future release. For now, you can't delete or tile any tabs; however, by right-clicking a tab and choosing Excluded, you'll place a strike through the name, and it'll be ignored.

8 QUICK COMMANDS

Another one of Vivaldi's strengths is its support for Quick Commands—press Ctrl + E to bring it into focus. Think of it as Vivaldi's alternative to Launchy (www.launchy.net), allowing you to bypass menus and mouse clicks in favor of quick keystrokes. Type a keyword, and it'll bring up related search terms covering everything from the usual web search results and existing workspaces and tabs, to commands, notes, extensions, and more. Visit Settings > Quick Commands to choose search types, and what priority should be given to displaying each type within the results.

» Scroll down, and you'll come to a section marked Command Chains, which unlocks a macro-like tool that runs multiple commands from the Quick Commands bar or via a keystroke or mouse gesture. Three examples are supplied—the 'Open Links and Tile' tab is best examined to see how you can combine commands to automatically open three links in their own tabs, before grouping them together in a single stack and tiling them vertically. Here, it's simple to create (and test) your own Command Chains before switching to the Keyboard section of Settings and scrolling down to Chains, where you can assign a keyboard shortcut to your chain if required. ⏻

Secure your Secrets with VeraCrypt

YOU'LL NEED THIS

VERACRYPT
https://veracrypt.fr, available for Windows, macOS and Linux

IT'S HARD TO GET ENCRYPTION RIGHT. Criminals using the supposedly end-to-end encrypted app Encrochat discovered this the hard way in 2020 when law enforcement was given permission to infect the service's servers with malware, allowing messages to be read.

VeraCrypt seeks to solve the problem of how to encrypt data at risk. It was originally developed in 2013 as a fork of the discontinued 'Truecrypt' project. Since then, it has made considerable improvements on Truecrypt's original labyrinthine and occasionally buggy code.

For Windows users, the software can encrypt your OS and external disks. It can also be used to create encrypted 'volumes'—ordinary files that can be placed on your device. When mounted as a virtual disk, VeraCrypt functions as an encrypted container for your private data.

VeraCrypt is open-source and also available for macOS and Linux, though there's no option to encrypt your OS for non-Windows system partitions. In this tutorial, we're focusing on the Windows version of VeraCrypt, available as both an MSI and EXE installer.

If you're considering VeraCrypt for Windows, disable Bitlocker protection via the Control Panel or PowerShell. Visit https://learn.microsoft.com/en-us/troubleshoot/windows-client/windows-security/suspend-bitlocker-protection-non-microsoft-updates. –**NATE DRAKE**

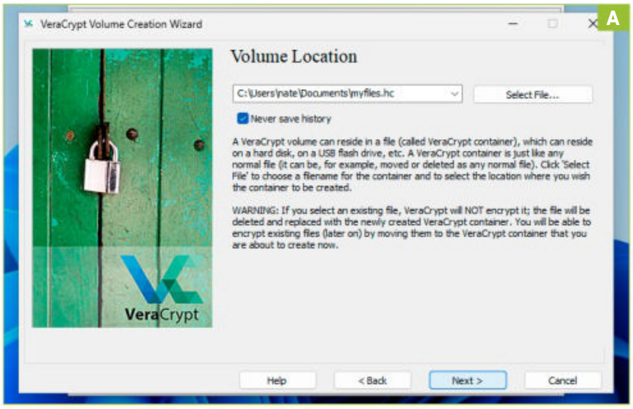

Volume Location

C:\Users\nate\Documents\myfiles.hc

☑ Never save history

A VeraCrypt volume can reside in a file (called VeraCrypt container), which can reside on a hard disk, on a USB flash drive, etc. A VeraCrypt container is just like any normal file (it can be, for example, moved or deleted as any normal file). Click 'Select File' to choose a filename for the container and to select the location where you wish the container to be created.

WARNING: If you select an existing file, VeraCrypt will NOT encrypt it; the file will be deleted and replaced with the newly created VeraCrypt container. You will be able to encrypt existing files (later on) by moving them to the VeraCrypt container that you are about to create now.

1 YOUR FIRST VOLUME

Click 'Create Volume' to start the volume creation wizard. At first, we suggest choosing to create an 'encrypted file container' in 'standard' format. Choose 'Select File' to choose where to place it and use the '.hc' extension to ensure that it will be opened by VeraCrypt.

» The 'encryption options' defaults to AES and SHA-512. Use the 'test' and 'benchmark' buttons to test how quickly this and other combinations of cipher and hash algorithms will encrypt/decrypt on your system.

» Use 'Volume Size' to determine the container size. VeraCrypt also supports creating dynamic sparse-file containers, though these won't perform nearly as efficiently as those of a fixed size. They also can't contain hidden volumes. Still, if you do want a dynamic container, use this window to set its maximum size.

» With respect to VeraCrypt's advice on choosing a passphrase with special characters and numbers, we recommend using Diceware (https://diceware.dmuth.org) to generate six or more random dictionary words for this purpose.

» You can also check 'use keyfiles' here and click the corresponding button to select any number of keyfiles to combine with the passphrase. They can be selected in any order, but make sure you read the red warning message. VeraCrypt itself will never modify keyfiles in any way, but if any part of the first 1024 KB of any keyfile changes, you won't be able to mount the encrypted volume with it.

» Use the 'Volume Format' window to specify a filesystem. You can reformat a mounted container to use a different filesystem at a later date, if you wish. From here, you can also choose to create a 'Dynamic' container. Move your mouse around to increase the available entropy for volume creation. [Image A]

2 MOUNT/DISMOUNT YOUR VOLUME

After exiting the wizard, VeraCrypt returns to the main screen. Choose 'Select File' and navigate to your chosen volume. Next, click a vacant drive letter in the main window

» You can now click 'Mount' at the bottom left to attempt to access the volume as a virtual drive.

» The password screen is self-explanatory. Enter your passphrase and click 'Keyfiles' to add one or more. Again, you can add keyfiles in any order. After you click 'OK', VeraCrypt will mount the volume using the first available drive letter.

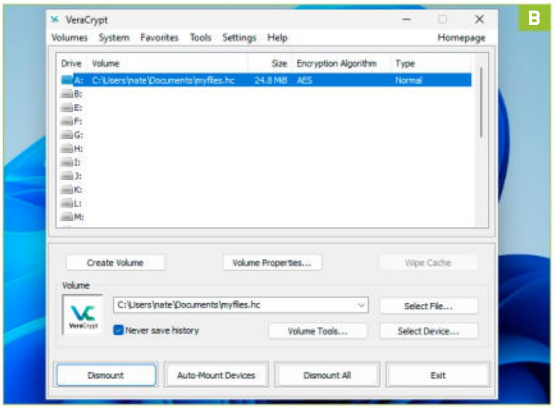

Drive	Volume	Size	Encryption Algorithm	Type
A:	C:\Users\nate\Documents\myfiles.hc	24.8 MB	AES	Normal
B:				
E:				
F:				
G:				
H:				
I:				
J:				
K:				
L:				
M:				

C:\Users\nate\Documents\myfiles.hc

☑ Never save history

© VERACRYPT

» By default, VeraCrypt won't remember your history, but click 'Favorites' on the menu bar in the main window to add your chosen volume for quick access next time.

» If you need to change the passphrase for your chosen volume, click on the 'Volumes' menu > 'Change Volume Password'. From here, you can also select 'Add/Remove keyfiles to/from Volume'.

» If mounting is successful, you can double-click on the drive letter to see the volume mounted as a virtual disk in File Explorer.

» You can now copy your private files onto the disk to place them inside the volume. If these files already exist on your hard drive, we recommend using a dedicated utility like Eraser (https://eraser.heidi.ie) to overwrite the original files with random data.

» Once all file operations are complete, return to the VeraCrypt main window. Click 'Dismount' to unmount the virtual disk. If you need more space for your chosen volume, choose 'Tools' from the main menu, then 'Volume Expander' to increase its size.

» Once an encrypted container or drive is dismounted, you can perform more operations on it via the 'Volume Tools' button. Aside from offering an easy way to change your password and manage keyfiles, you can also choose to 'Backup Volume Header'. This means that if the encrypted volume header containing the master key and salt become damaged, you can restore your backup copy and access your files. [**Image B**]

3 CREATE HIDDEN VOLUMES

Setting up a hidden volume works similarly to creating a standard one. In the 'Volume Creation Wizard'—just choose 'Hidden VeraCrypt Volume' in the 'Volume Type' window.

» 'Normal' mode creates a standard VeraCrypt volume with a hidden one alongside it, and 'Direct' mode allows you to create a hidden volume inside an existing VeraCrypt volume. If you use this mode, dismount your VeraCrypt volume/drive first.

» Assuming you choose 'Normal' mode, VeraCrypt will ask you to configure the 'outer' volume, which is designed to be discovered by an adversary. This offers the same options as for creating a standard volume, such as setting your own encryption algorithm, passphrase, and keyfiles. Once you've set these and chosen the filesystem, VeraCrypt will format the outer volume and fill it with random data.

» Once this is done, the outer volume will automatically be mounted. Click 'Open Outer Volume' to launch File Explorer and copy over some harmless files you don't mind being discovered.

» Next, the outer volume will be dismounted and the Hidden Volume Wizard will launch. VeraCrypt will automatically determine the maximum possible size of the hidden volume. Naturally, this will always be smaller than the outer volume.

» The hidden volume operates independently of the outer one in that you can choose a different encryption cipher, hash algorithm, and filesystem. You should also set a strong passphrase and (optionally) keyfiles. Naturally, this should be different to the passphrase/keyfiles you use for the outer volume.

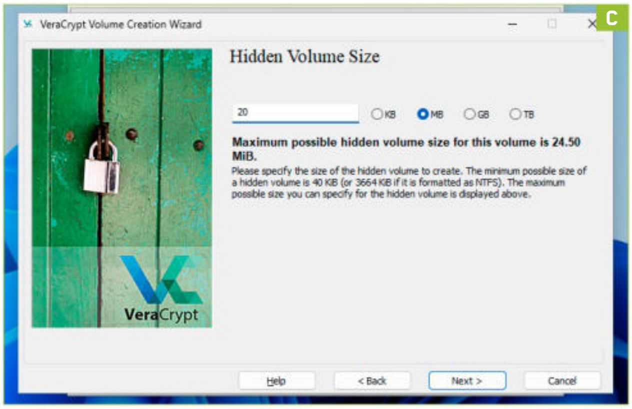

VeraCrypt Volume Creation Wizard

Hidden Volume Size

20 ○ KB ● MB ○ GB ○ TB

Maximum possible hidden volume size for this volume is 24.50 MiB.

Please specify the size of the hidden volume to create. The minimum possible size of a hidden volume is 40 KiB (or 3664 KiB if it is formatted as NTFS). The maximum possible size you can specify for the hidden volume is displayed above.

C

» In future, if an adversary forces you to reveal the password for your VeraCrypt volume, input the 'outer' password to show them the dummy files you placed there earlier. The hidden volume will simply present as free space filled with random data. [**Image C**]

4 MOUNT HIDDEN VOLUMES SAFELY

If you mount a VeraCrypt volume containing a hidden volume with the 'outer' password, you may read data stored on the outer volume without any risk.

» However, as we've learned, VeraCrypt's security is double blind. Until mounted successfully, the software has no way to tell if an encrypted container or drive has a hidden volume. If only the outer volume password is entered, the hidden volume appears as free space.

» This feature is crucial to the security of hidden volumes. Still, it means that if any new data is moved to a mounted outer volume, there is a risk that the hidden volume header and/or its data will be overwritten.

» If you do need to copy more data to the outer volume, there's an easy way to prevent this. First, use 'Select File' or 'Select Device' to choose your VeraCrypt Volume, then choose 'Mount'.

» On the password window, click 'Mount Options', then check the box under 'Hidden Volume Protection'. Here, you can enter the passphrase and specify keyfiles (if any) for the hidden volume. Click 'OK' to return to the main dialog and enter the passphrase for the outer volume.

» Once mounted, VeraCrypt will allow you to write more data to it safely. If you try to write any data to the hidden volume, VeraCrypt will attempt to set it to

CAVEAT ENCRYPTOR

The VeraCrypt website is quite clear about situations in which the program's encryption can't help you. If, for instance, you unlock an encrypted system partition containing malware, your data is still at risk.

If an adversary has physical access to your device, it could also be subject to an 'Evil Maid Attack', such as installing a hardware keylogger.

You can disguise the fact that you're using VeraCrypt system encryption to some extent via the main menu (Settings > System Encryption > Edit Boot Loader Configuration). From here, you can change the message that displays on first boot—for example, to 'Missing operating system'. You can find examples of this from https://github.com/VeraCrypt/VeraCrypt/issues/758 . If your device was seized and examined, though, it would be simple for an adversary to spot the Boot Loader.

Plausible deniability is also a moot point if you create VeraCrypt containers—even if you remove the telltale .hc extension, it would be very hard to explain why you have files of seemingly random data on your machine.

Above all, remember that when creating a new VeraCrypt Volume, if you choose an existing file, it will be replaced with the VeraCrypt container. If you choose to encrypt a partition or external drive instead, any existing data on it will be lost during formatting.

TOUR VERACRYPT'S MAIN INTERFACE

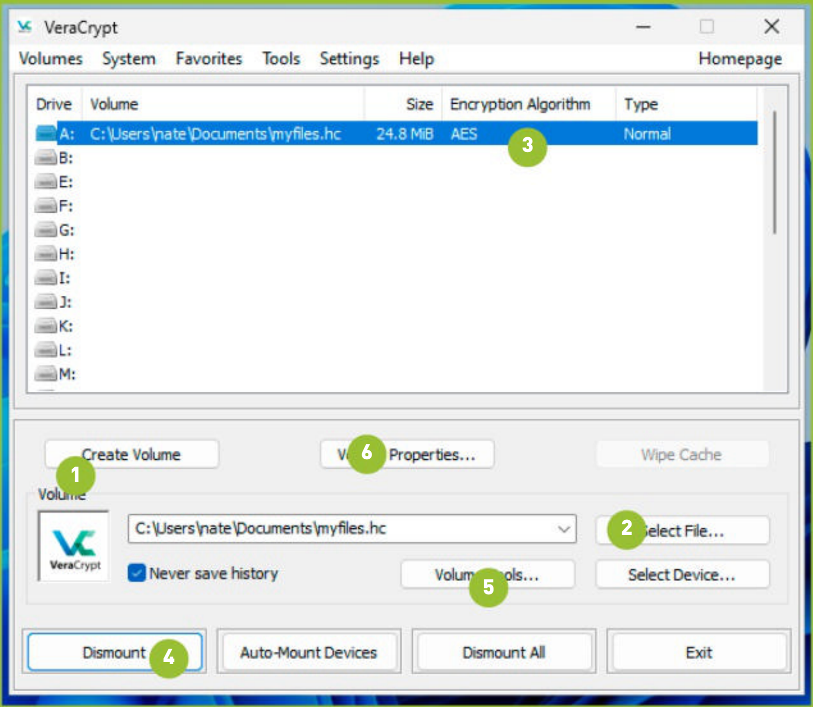

Drive	Volume	Size	Encryption Algorithm	Type
A:	C:\Users\nate\Documents\myfiles.hc	24.8 MiB	AES	Normal
B:				
E:				
F:				
G:				
H:				
I:				
J:				
K:				
L:				
M:				

Volume: C:\Users\nate\Documents\myfiles.hc

☑ Never save history

1. VOLUME CREATION
Click to launch the volume creation wizard. This is used for creating VeraCrypt encrypted volume files, as well as external drives. The wizard also supports configuring key files and creation of hidden volumes.

2. SELECT FILE
If you've created an encrypted volume (.hc format), use 'Select File' to point to it in explorer. You can also use 'Select Device' to choose entire external drives or partitions encrypted via VeraCrypt.

3. VIRTUAL DRIVES
Once the correct passphrase is entered, mounted volumes/drives are accessed as virtual drives. On Windows, each is assigned its own letter and can be accessed via Explorer just as with an external disk.

4. MOUNT/DISMOUNT
Once all file operations complete, click 'Dismount' to unmount the virtual drive and erase the passphrase from RAM. If no virtual drives are selected, this button can be used to mount VeraCrypt volumes/disks.

5. VOLUME TOOLS
From here, you can change the volume password, add/remove keyfiles, set the header key derivation algorithm, as well as back up/restore the volume header. Volumes must be dismounted before carrying out most operations.

6. VOLUME PROPERTIES
If you highlight a mounted volume in the main window, then select 'Volume Properties', and you can choose to view a rundown of key information. This includes volume size, location, and encryption algorithm, in addition to key and block size.

'read only', so make sure to perform file operations only on the outer volume itself. In other words, if you want to write to the hidden volume specifically, be sure to mount the drive using the corresponding password for the hidden volume only.

» If you're unsure if you've set up safe mounting of the outer volume correctly, check the 'Type' category in the main VeraCrypt window, which should show 'Outer'. Click 'Volume Properties' to ensure that 'Hidden Volume Protected' is set to 'Yes'. [Image D]

5 STANDARD SYSTEM ENCRYPTION

Depending on the version of Windows you're running, VeraCrypt can encrypt the system partition/drive.

» Pre-boot authentication is handled by the VeraCrypt Boot Loader, which resides in the first track of the boot drive. It doesn't make use of TPM. Once system encryption is in place, any data written to the Windows partition is encrypted in memory, then written to the disk.

» To get started, just click into the VeraCrypt 'System' menu, then 'Encrypt System Partition/Drive'.

» The setup wizard allows you to specify 'Normal' or 'Hidden' system encryption. It also supports support multi-boot environments, e.g. Windows and Linux, though system encryption only applies to Windows partitions.

» Users can choose the same encryption/hash algorithms as when creating regular VeraCrypt, but using cascade ciphers here will reduce system performance, so we recommend staying with the default (AES). VeraCrypt supports parallelization, as well as compatible hardware for AES encryption, so choosing AES alone will likely result in the fastest speeds.

» The wizard also requires users to create a 'Rescue Disk' image (in ZIP format) and save it to an external drive. This is insurance against the boot loader or windows partition becoming corrupted. It contains both a backup of the VeraCrypt Boot Loader with encryption key, as well as the current EFI boot loader.

» VeraCrypt will also perform a pretest to make sure your system is compatible. Be sure to read the warning notes, then restart your PC to complete prechecks.

» If you're using a portable device, make sure it's plugged in to the mains before proceeding. In theory, you can pause/defer the encryption process at any time, and even shut down/restart your machine. However, we recommend closing any open files and applications and backing up your data before proceeding, as there's always a risk of data loss. [Image E]

6 CONFIGURE A 'HIDDEN' OPERATING SYSTEM

VeraCrypt system encryption offers a 'hidden operating system' alongside an 'outer' system onto which canny users can place some harmless-looking data and applications.

© VERACRYPT

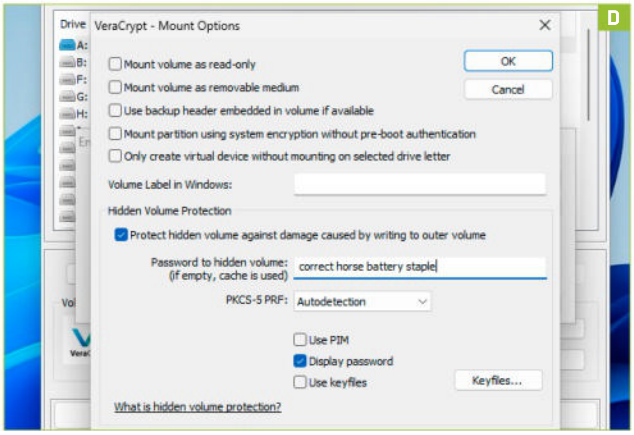
D

VeraCrypt - Mount Options

☐ Mount volume as read-only
☐ Mount volume as removable medium
☐ Use backup header embedded in volume if available
☐ Mount partition using system encryption without pre-boot authentication
☐ Only create virtual device without mounting on selected drive letter

Volume Label in Windows:

Hidden Volume Protection

☑ Protect hidden volume against damage caused by writing to outer volume

Password to hidden volume: (if empty, cache is used): correct horse battery staple

PKCS-5 PRF: Autodetection

☐ Use PIM
☑ Display password
☐ Use keyfiles

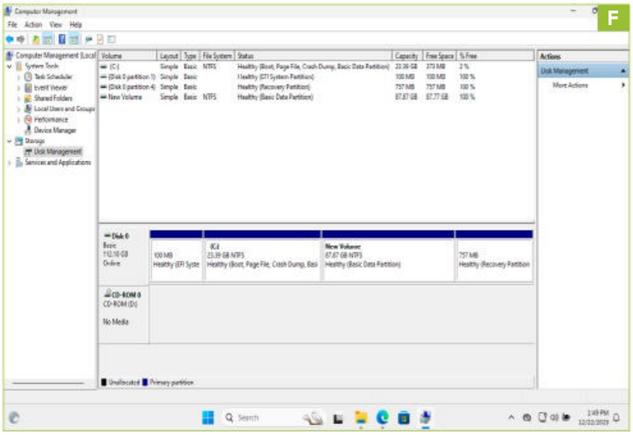
F

» During our tests, however, we weren't able to set this up on our Windows 11 test machine, despite creating a second partition for the hidden OS that was 110 percent larger than the system partition as specified in the documentation. We used a Gparted 'Live' CD to do this, though this is also possible using regular Windows disk management tools.

» Visit https://VeraCrypt.eu/en/VeraCrypt%20Hidden%20Operating%20System.html to discover more. We recommend making a full backup of your system before proceeding.

» The very fact that the 'hidden' option necessitates creating a second, encrypted VeraCrypt partition, and that the VeraCrypt Boot Loader is in plain sight, also means that this feature's 'plausible deniability' is called into question.

» You can mitigate this slightly by creating a VeraCrypt Rescue Disk with a Boot Loader, then removing the Boot Loader from the hard drive. Anyone analyzing the system partition will see that it appears only to contain random data, though naturally the presence of the rescue disk would still prove you're using VeraCrypt.

» As with standard system encryption, we recommend making a full system backup before you begin partitioning and encrypting drives. [**Image F**]

7

USE THE KEYFILE GENERATOR

Although using data already on your drive provides the best plausible deniability for keyfiles, you're running the risk that Windows will modify them. This is a particular risk for media files, as various players can update the metadata to reflect the title, artist, and so on.

» Using files which others can access, like your MP3 collection, also undermines the security of your VeraCrypt volume, given that the first 1MB of such files are non-random. Luckily, VeraCrypt comes to the rescue with its own Keyfile Generator.

» Before getting started, go to 'Settings > Performance/Driver Configuration'. Check the box marked 'Use CPU hardware random number generator as an additional source of entropy'. This will allow you to generate better-quality randomness rather than just relying on your mouse movements. You'll need admin privileges in order to enable this feature.

» Next, go to 'Tools > Keyfile Generator'. Leave the default PRF (SHA512) as is. Make sure to move your mouse around as randomly as possible until the progress bar is green.

» From here, you can also specify the number of keyfiles to generate, their size (default is 64 KB), and the base name, eg. mykeyfile.

» Naturally, the presence of keyfiles on your system does little for plausible deniability. Still, they can give you some peace of mind if, for instance, you upload your VeraCrypt container to a cloud storage service like Dropbox.

» Provided the keyfiles are kept offline, you can rest assured that no one will be able to access the container, even if your password is leaked.

» If you want to throw off adversaries, consider generating a large number of numbered keyfiles, but only use certain ones, like 3, 7 and 11, to unlock your VeraCrypt volume. VeraCrypt doesn't log keyfiles you use, nor alter them, so there would be no way to tell the real keyfiles from the redundant ones. [**Image G**]

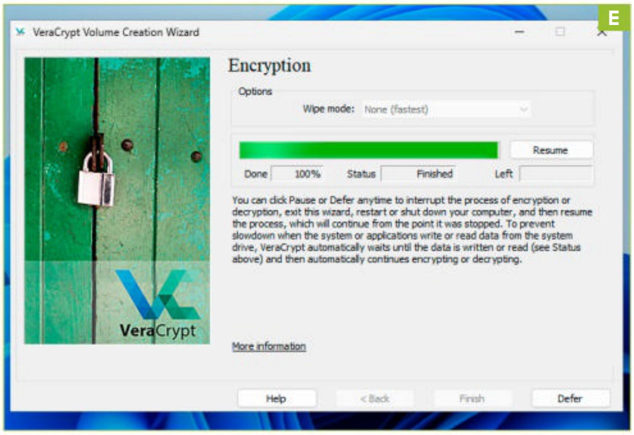
E

VeraCrypt Volume Creation Wizard

Encryption

Options

Wipe mode: None (fastest)

Done: 100%
Status: Finished
Left:

You can click Pause or Defer anytime to interrupt the process of encryption or decryption, exit this wizard, restart or shut down your computer, and then resume the process, which will continue from the point it was stopped. To prevent slowdown when the system or applications write or read data from the system drive, VeraCrypt automatically waits until the data is written or read (see Status above) and then automatically continues encrypting or decrypting.

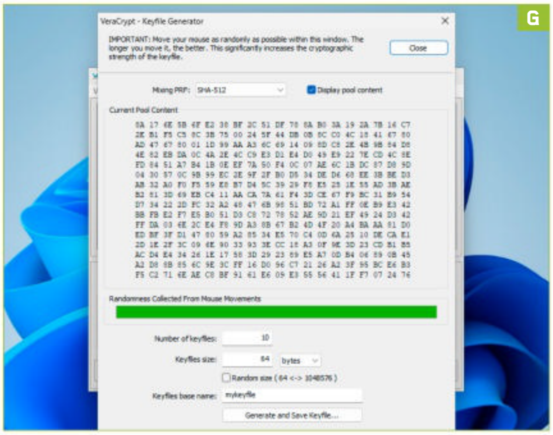
G

VeraCrypt - Keyfile Generator

IMPORTANT: Move your mouse as randomly as possible within this window. The longer you move it, the better. This significantly increases the cryptographic strength of the keyfile.

Mixing PRF: SHA-512
☑ Display pool content

Current Pool Content

Randomness Collected From Mouse Movements

Number of keyfiles: 10
Keyfiles size: 64 bytes
☐ Random size (64 <-> 1048576)
Keyfiles base name: mykeyfile

Use Bitwarden to store your passkeys

YOU'LL NEED THIS

BITWARDEN
www.bitwarden.com

PASSKEYS, which are set to largely replace passwords over the coming years, are strings of code synchronized between your devices that are used to prove who you are. As we've covered previously, they're more secure than passwords and easier to use, because you don't need to remember them. Traditionally, they were stored on specific devices like a phone or security key, which you'd need to keep with you to log into a website. However, password managers—like Bitwarden—are now adding support for passkeys. Here, we explain how to use Bitwarden to log into websites without using passwords. –**NIK RAWLINSON**

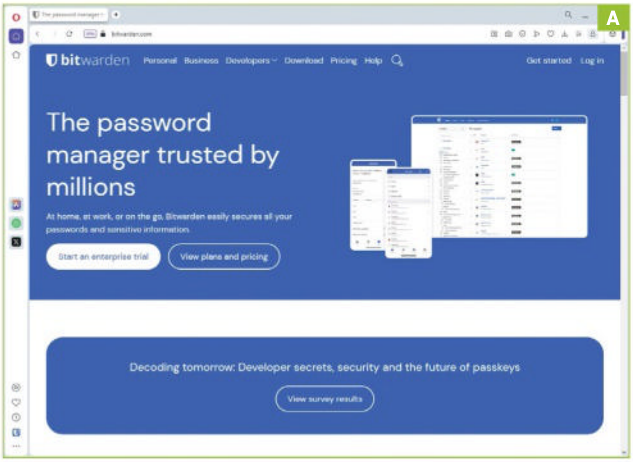
A
bitwarden
Personal Business Developers Download Pricing Help
Get started Log in
Start an enterprise trial
View plans and pricing
View survey results

The password manager trusted by millions

At home, at work, or on the go, Bitwarden easily secures all your passwords and sensitive information.

Decoding tomorrow: Developer secrets, security and the future of passkeys

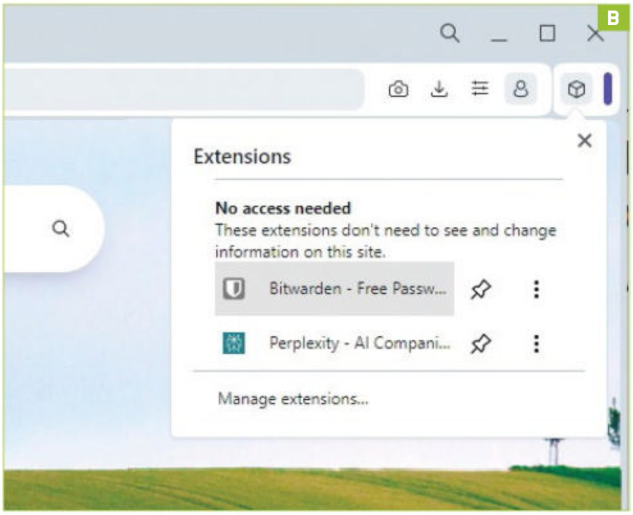
B
Extensions
No access needed
Bitwarden - Free Passw...
Perplexity - AI Compani...
Manage extensions...

These extensions don't need to see and change information on this site.

1 SET UP AN ACCOUNT

If you don't already have a Bitwarden account, set one up by visiting www.bitwarden.com and clicking 'Get started' [**Image A**]. You'll need to provide your email address (for signing in to your Bitwarden account), name, and master password. The master password is used to lock all of your logins, so it's very important that you don't forget it. You can enter a master password hint to jog your memory, but don't make this so obvious that it would help someone else access your account.

2 INSTALL THE EXTENSION

Now, download the Bitwarden extension for your browser. Although Bitwarden also has a downloadable program, passkey support is initially being rolled out only to its browser extensions. Go to www.bitwarden.com/download [**Image B**], scroll down to Web Browser Extensions, and click your browser in the list to be taken to its add-on page. Click the install button (or Get if you're using Edge). If it doesn't automatically offer to log you in when it's finished installing, select Bitwarden in your browser's extensions menu, and then log in using the email address and password you specified in Step 1.

3 CHOOSE A WEBSITE OR SERVICE

The process for creating a passkey varies between services. Here, we'll set one up to log into a Microsoft account. Start by opening a new browser window and visit www.office.com [**Image C**]. Log in using your Microsoft account details. Click your profile icon at the top right of the window, then click 'My Microsoft account'. Click Security and re-enter your password if requested. Click 'Advanced security options', followed by 'Add a new way to sign in or verify'.

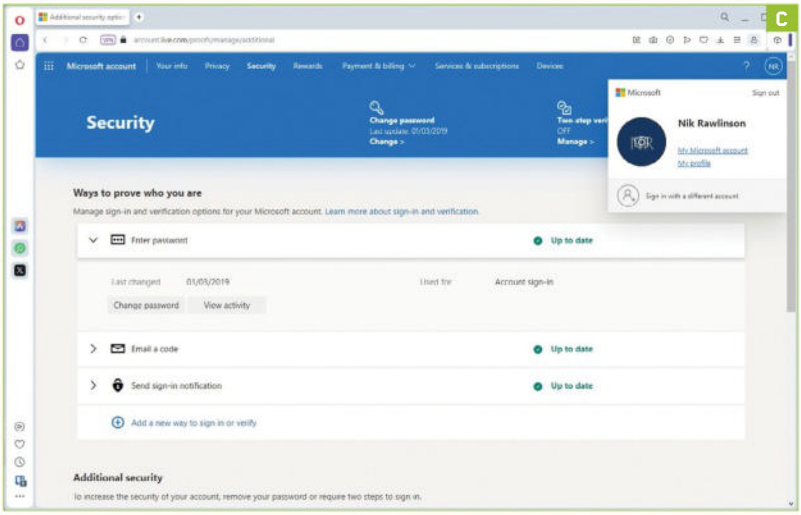
C
Microsoft account
Your info
Privacy
Security
Rewards
Payment & billing
Services & subscriptions
Devices
Security
Change password
Last update: 01/03/2019
Change >
Two-step veri
OFF
Manage >
Microsoft
Sign out
Nik Rawlinson
My Microsoft account
My profile
Sign in with a different account
Enter password
Up to date
Last changed
01/03/2019
Used for
Account sign-in
Change password
View activity
Email a code
Up to date
Send sign-in notification
Up to date
Add a new way to sign in or verify

Ways to prove who you are

Manage sign-in and verification options for your Microsoft account. Learn more about sign-in and verification.

Additional security

To increase the security of your account, remove your password or require two steps to sign in.

© BITWARDEN

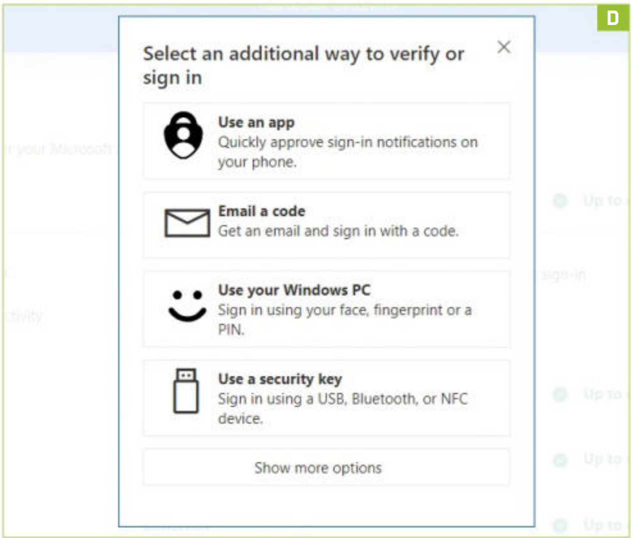

Select an additional way to verify or sign in

Use an app
Quickly approve sign-in notifications on your phone.

Email a code
Get an email and sign in with a code.

Use your Windows PC
Sign in using your face, fingerprint or a PIN.

Use a security key
Sign in using a USB, Bluetooth, or NFC device.

4 CREATE A PASSKEY

Click 'Use a security key' [**Image D**]. Traditionally, this is what you would have selected if you had a physical security key device, like a YubiKey, which is a USB dongle that plugs into your PC and is used for two-step authentication. Here, Bitwarden lets you use a synced software key instead of a physical device to do the same job more conveniently. Leave 'USB device' selected on the next page, and click Next. Bitwarden will pop up an alert asking you to 'Choose a login to save this passkey to'. Click 'Save passkey' and—when prompted to give it a name—type Bitwarden.

5 CHECK YOUR PASSKEY

Before going any further, you should check that your passkey works. Log out of your Microsoft account by clicking your profile icon in the top-right corner, then click 'Sign out'. You'll return to the sign-in window, which will prompt you for your email address. Don't enter it. Instead, click 'Sign in with Windows Hello or a security key'. The screen will refresh and a Bitwarden window will pop up, asking 'Log in with passkey?' [**Image E**]. Check that the address shown in the window is the same as the one in the address box before the first slash. If so, click Confirm.

6 REMOVING A PASSKEY

You will now be logged into your Microsoft account without having needed to enter a password—and, because the passkey is saved in your Bitwarden account, not just in your browser; you'll be able to do the same using any other browser running the Bitwarden extension. If you need to delete the passkey from your Microsoft account (perhaps because you've lost access to your Bitwarden account), return to the Security page (see Step 3), and click the arrow next to 'Use a security key' [**Image F**] to expand that section. Now, click the 'Remove' button below your key.

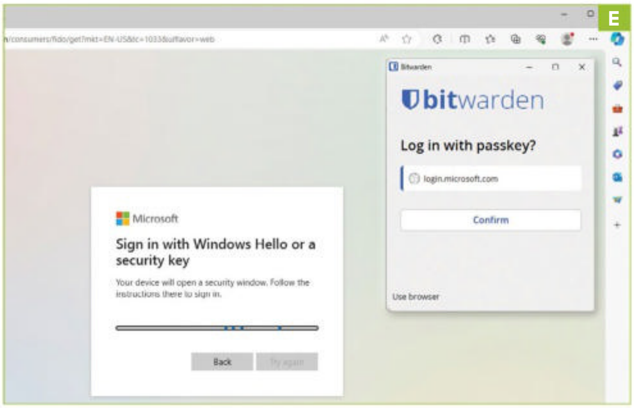

Bitwarden

bitwarden

Log in with passkey?

login.microsoft.com

Microsoft

Sign in with Windows Hello or a security key

Your device will open a security window. Follow the instructions there to sign in.

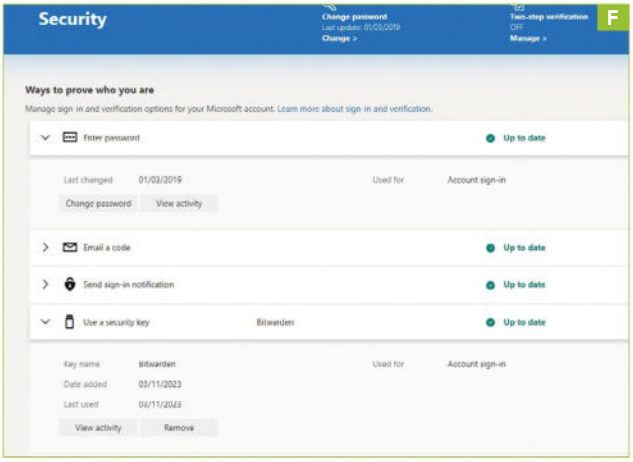

Security

Change password
Last update: 01/03/2019
Change >

Two-step verification
OFF
Manage >

Ways to prove who you are

Manage sign in and verification options for your Microsoft account. Learn more about sign in and verification.

Enter password — Up to date

Last changed 01/03/2019 — Used for Account sign-in

Email a code — Up to date

Send sign-in notification — Up to date

Use a security key — Bitwarden — Up to date

Key name Bitwarden — Used for Account sign-in

Date added 03/11/2023

Last used 03/11/2023

7 DELETING A PASSKEY

If you no longer want to use Bitwarden to access a website or service, it's good practice to not just remove the passkey from the site in question, but to also delete it entirely from Bitwarden. To delete a key from your Bitwarden account, navigate to your browser's Extensions menu, and then select Bitwarden from the list. Log in if it asks for your Bitwarden password, then click 'Vault' at the bottom of the box. After this, choose your Microsoft login (or whichever key you want to delete) to open its details, and click 'Delete item' [**Image G**] at the bottom of the panel. ⏻

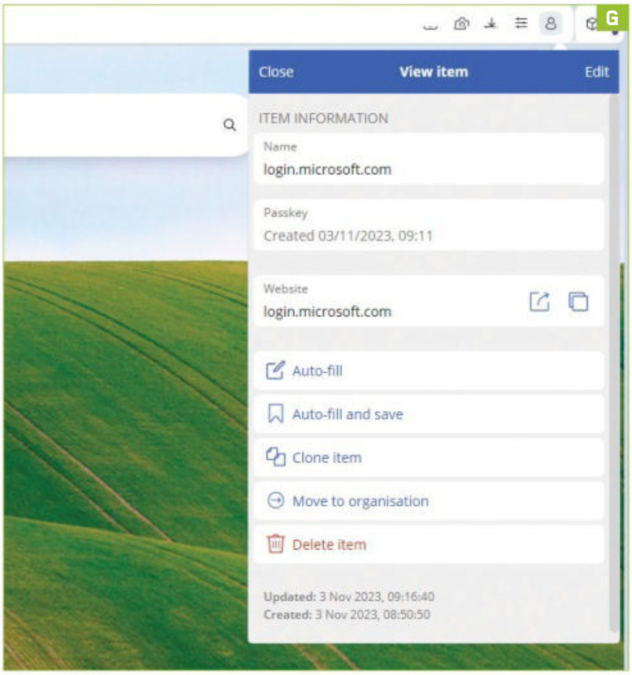

ITEM INFORMATION

Name
login.microsoft.com

Passkey
Created 03/11/2023, 09:11

Website
login.microsoft.com

Updated: 3 Nov 2023, 09:16:40
Created: 3 Nov 2023, 08:50:50

Share files from PC to phone for free

YOU'LL NEED THIS

LOCALSEND
www.localsend.org
For Windows, plus iOS
or Android

CLOUD SERVICES, like Google Drive and iCloud, simplify the task of sharing files between your phone and computer. The trouble is that when you share this way, it's easy to end up with loose files cluttering your cloud storage once you've completed the transfer.

This wastes disk space, so unless you tidy up after yourself, you may one day need to pay for additional storage. Alternatively, you could use LocalSend, which rather than routing files via the cloud, sends them directly across your network. It's free, with no restrictions, and works a lot like Apple's AirDrop, without requiring all your devices to be made by Apple. –**NIK RAWLINSON**

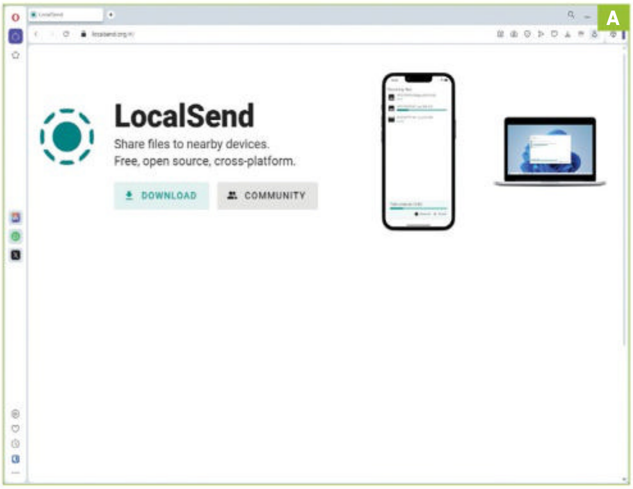

LocalSend
Share files to nearby devices.
Free, open source, cross-platform.

1 DOWNLOAD AND INSTALL

Open www.localsend.org and click Download [**Image A**]. Click 'Windows' on the following screen, followed by 'MSIX'. This downloads a version of LocalSend that you can install. Launch the downloaded file and click Install. If you prefer a version that doesn't need installing before you launch it, click 'Zip (Portable)', extract the compressed files, then launch 'localsend_app. exe'. If a box appears warning that 'Windows protected your PC', click 'More info' followed by 'Run anyway' to launch LocalSend.

2 GET LOCALSEND FOR YOUR SMARTPHONE

Download the free LocalSend app for Android (play.google.com/store/apps/details?id=org.localsend.localsend_app) or iOS (apps.apple.com/us/app/localsend/id1661733229). If you have an iPhone, it will ask for permission to find devices on your local network. In our tests, we've found that while refusing this doesn't stop an iPhone receiving files from a PC, it does prevent it from sending them. So, tap 'Allow' [**Image B**] when the dialog box pops up. Android doesn't ask this question, so you can go straight to the next step.

3 GET CONNECTED

With both your phone and your computer connected to the same wireless network, you can now send files between them without needing usernames, passwords, or any intermediate service like Dropbox. It's a lot like Apple's AirDrop, but without being locked into Apple's walled garden. Notice how every device that you've connected has been given its own unique name. As you can see from the photo [**Image C**], our iPhone, Android phone and PC each have a fruit- or vegetable-based name by which we can identify them.

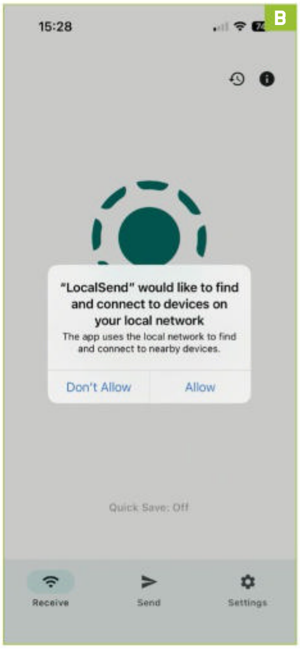

"LocalSend" would like to find and connect to devices on your local network
The app uses the local network to find and connect to nearby devices.

C
Fantastic Pumpkin
Great Broccoli
#84 #122
Fast Blueberry
#14

© LOCALSEND

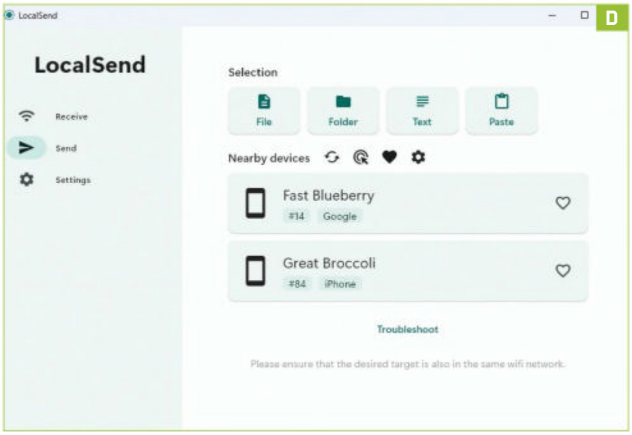
D
LocalSend
LocalSend
Receive
Send
Settings
Selection
File
Folder
Text
Paste
Nearby devices
Fast Blueberry
#14
Google
Great Broccoli
#84
iPhone
Troubleshoot

Please ensure that the desired target is also in the same wifi network.

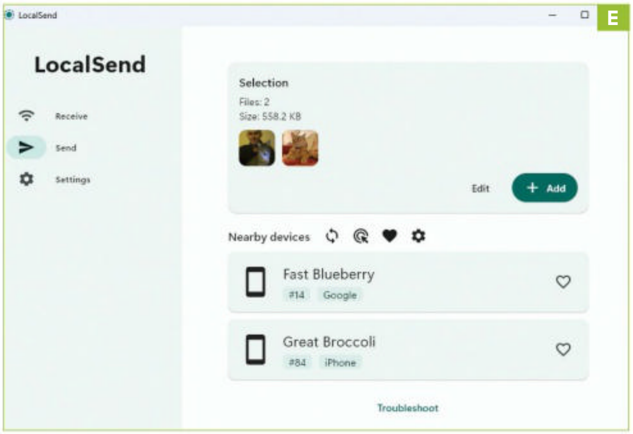
E
LocalSend
LocalSend
Receive
Send
Settings
Selection
Edit
+ Add
Nearby devices
Fast Blueberry
#14
Google
Great Broccoli
#84
iPhone
Troubleshoot

Files: 2
Size: 558.2 KB

4 SEND FILES FROM PC TO PHONE

To send a file from your PC to your phone, click Send, then click the closest match in the Selection section to what you want to send. We're going to send a photo from our computer to our iPhone, so we'll click File [**Image D**]. Navigate to the photo in the File Explorer window that appears, select it, and click Open. The file will replace the original contents of the Selection panel, where a new 'Add' button will appear. Click this, and repeat if you want to send several files simultaneously.

5 CHOOSE TARGET DEVICE

Click the device to which you want to send the selected file. Here, because we want to send the file to our iPhone, we'll click Great Broccoli [**Image E**]. We know this is correct because it's marked as an iPhone in the 'Nearby devices' panel, while the Android device is labeled Google. If we had multiple iPhones or multiple Android devices, you can make sure you're sending files to the right one by matching the name in this panel with the one displayed on the screen of the phone itself.

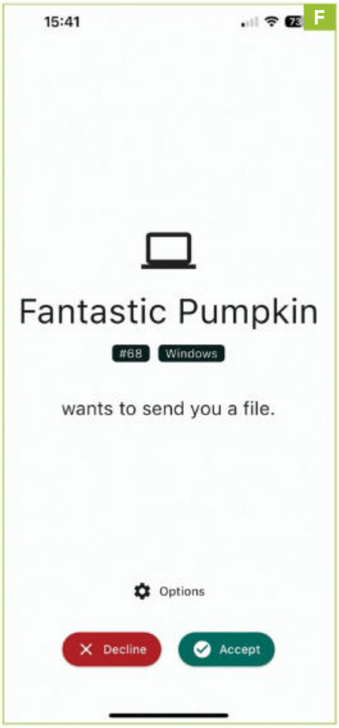
F
15:41
Fantastic Pumpkin
#68
Windows
Options
Decline
Accept

wants to send you a file.

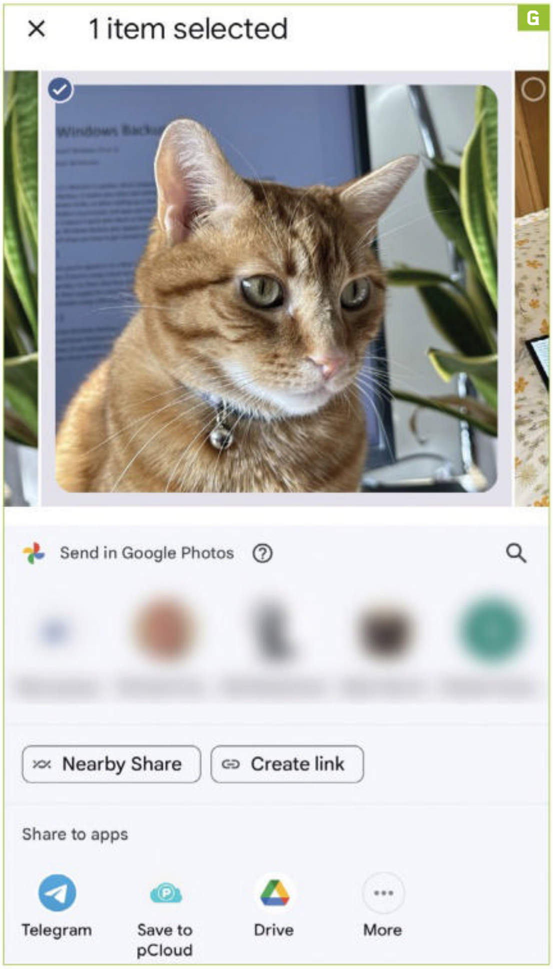
G
1 item selected
Send in Google Photos
Nearby Share
Create link
Share to apps
Telegram
Save to pCloud
Drive
More

6 ACCEPT THE TRANSFER

Check your phone screen, where you should see an alert that your computer wants to send you a file. As long as the name on the screen [**Image F**] is the same as the unique name of your computer (so you know the file is coming from your own PC and not someone else using LocalSend), tap Accept. While our Android phone automatically saved the transferred photo to our Photos library, our iPhone asked permission to add it to Photos. If yours does this, tap OK to accept the transfer.

7 BUMP FILES BACK TO YOUR PC

The process works the same in the opposite direction. To send a file from an iPhone to your PC, open it, then tap the share button (a square with an arrow pointing up out of it). Scroll the second row of icons, tap LocalSend, then select your PC's fruit or veg name as the destination. On Android, open the file you want to send, and tap Share (it looks like a left-pointing chevron). Tap More on the 'Share to apps' line [**Image G**], then tap LocalSend, and tap the name of your PC. The files are saved to your PC's Downloads folder. ⏻

LAB NOTES

JEREMY LAIRD, CONTRIBUTOR

OLED is on the up

So nearly awesome

You can already buy new 32-inch 4K OLED monitors, like the Alienware 32 AW3225QF.

© ALIENWARE

WHEN IT COMES to PC technology, this year's CES in Vegas was dominated by new OLED screens. Depending on how you measure it, we're on to the second or third generation of OLED display technology in gaming monitors, and the show was rammed with new models in all shapes and sizes.

Want higher refresh rates? Have some. More pixels? No problem. Improved full-screen brightness? Kinda. Arguably, it's a new generation of 32-inch 4K OLED monitors running at 240Hz that seem to tick the most boxes. Remarkably, these aren't prototype displays. By the time you read these words, you'll be able to jump on Dell's website and order the Alienware 32 AW3225QF for $1,199. It's real and it's here.

It uses Samsung's second-gen QD-OLED technology, produced via inkjet printing to enable the higher pixel density required to achieve 4K on a 32-inch panel. Previously, all 4K gaming monitors had been based on large TV panels of at least 42 inches. I can't wait to see what these high-DPI OLEDs look like.

On the speed side, several brands announced 360Hz OLED monitors, but Asus edged everyone out with a sole 480Hz 1440p panel, the ROG Swift PG27AQDP. It won't be available until later this year, but when it arrives, it will pretty much close the door on one of the few remaining areas where OLED technology lags behind LCD.

The other obvious measure by which OLED doesn't match up to LCD is full-screen brightness. LG's panels in particular have been poor by this metric, and 2024 sees the arrival of the company's second-gen MLA or micro lens array OLED panel tech, again to be used by numerous monitor brands.

It's a definite improvement, largely closing the gap to Samsung's QD-OLED panels. But it still only hits around 250 nits full screen. As it happens, Samsung's second-gen QD-OLED panels only go slightly brighter, so there's still work to be done. I think around 400 nits is required for full-screen brightness to be a non-issue.

The other problem is pricing. Most of them are well over $1,000. It's difficult to understand why a 32-inch 4K monitor is $1,200, but a 50-inch 4K OLED TV can be had for similar money. 34-inch ultrawide panels look awfully pricey, too.

With OLED looking set to get even better in the next few years, those prices are hard to stomach. As excited as I am, then, about OLED monitors, pulling a trigger on one today could be a recipe for regret in fairly short order.

ZAK STOREY

Contributor

Intel's absence from the handheld gaming PC market has been a bit glaring, so launching the MSI Claw using Intel's Meteor Lake is news.

Let's hope Meteor Lake isn't proving too hot in the Claw—that wouldn't be a good thing. Anyway, what is remarkable is how close Meteor Lake, branded Core Ultra, seems to be to AMD's Phoenix APU, also known as the Ryzen Z1 Extreme.

It's amazing to think that two engineering teams designed complex SoCs, with CPU cores, GPU cores, I/O, memory controllers, and even AI-accelerating NPUs, then had them built on different production nodes and ended up with almost the same performance. It's a close call when it comes to frame rates between Intel's New Core Ultra in the MSI Claw and existing handhelds with the AMD Ryzen Z1 extreme chip.

Meteor Lake does seem to be competitive, which has to be a PR win. That said, I suspect that when the dust settles, AMD-powered handhelds will remain the pick for the most part.

But as much as Intel has massively improved its Arc graphics drivers over the past year or so—and it really has—I still think driver consistency with Intel graphics is a worry. Ultimately, it's hard to avoid the conclusion that you're likely to get a more consistent experience with an AMD handheld than one powered by Intel's new chip.

Reviewed…

74 Intel Core i7-14700K

76 Steam Deck OLED
78 RCA Evolution Premium M27PG135F

80 Hyte Y70 Touch
82 MSI MEG Z790 ACE MAX
84 Asus ZenScreen Ink MB14AHD
87 32GB XPG Lancer Blade RGB DDR5
88 Sony H9 Wireless Noise Cancelling Gaming Headset
89 WD Black SN770M PCIe 4.0 M.2 SSD

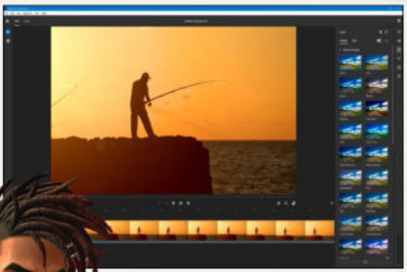

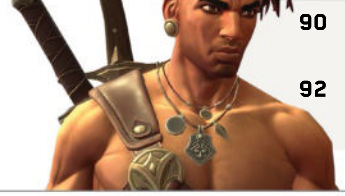

90 Prince of Persia: The Lost Crown
92 Microsoft Clipchamp vs Premiere Rush

Editor's Pick: Shure 55SH Series II

Finally getting some use

BIT OF A left field pick this, I'll admit, but it's one of the best things I've ever owned. For context, I come from a very musical family. The young 'uns, particularly big into metal and rock in general, are often gifted with guitar lessons, drum tuition, or any form of musical instrument training they desire as teenagers going through high school.

I spent an enormous amount of time in my youth playing in bands as a rhythm guitarist, and singing too (as I was one of the few in my cohort that could actually hold a note). That said, hitting 18, it has almost become a family tradition of ours to be given a piece of musical equipment of legendary status. For my uncle who's eight years older than me, it was a 1998 Cherry Sunburst Gibson Les Paul guitar, which he's still got to this day, and actually managed to get Billy Gibbons, of ZZ-Top fame, to sign it at a hot-rod show. For my cousin, he nabbed himself a Marhsall JCM2000 half stack setup. For me? Well, I turned 18 during the middle of the 2008 financial crash, and things were a little bit different.

All I wanted was something simple, you know? I knew my parents were struggling, but they still wanted to continue the tradition, so being a rhythm guitarist and singer, I asked for a classic: the Shure 55SH Series II Unidyne microphone. The antique, yet beautiful, die-cast, microphone of glory. Sadly, however, with college and tech journalism right around the corner, it wouldn't be until recently—nearly 12 years on—that I'd actually put it to good use.

How so? Simply put, it's now my daily driver for my streaming, voiceover IP, podcasting, and video-recording needs. Pretty much everyone I have calls with loves it, including my mother, who now lives in New Zealand. I've played around with a lot of dedicated XLR mics now, hooking them up to Elgato's Wave XLR interface, but nothing even comes close to this pure brilliance, and the absolute clarity and warmth that I can achieve with it. It's just ludicrously good. Impressively, it's still just $199, too.

It's got a shock-mounted cardioid cartridge, with integrated pop-filter and wind-breaker, along with a hefty mounting system, tilt functionality, and an integrated on-off switch. In short, it's an absolute king of dedicated XLR microphones.

Downsides? Well, it's heavy, that's a fact, and without tweaking it can be a little lacking on the bass (it's definitely more tailored to vocal audio than anything else), but other than that, the value on offer here, and the style, make it a surefire pick, certainly compared to the ProCast from HyperX, or Wave DX from Elgato. I haven't got it signed by James Hetfield yet, but you know, a boy can dream **–ZS**

$199, www.shure.com

SHURE
55SH SERIES II
OFF
UNIDYNE
DYNAMIC

It turns out that the Shure 55SH Series II was worth the wait.

© SHURE, UBISOFT, ADOBE

Intel Core i7-14700K

Is it a Core i9 or a Core i7?

INTEL CORE
i7-14700K
SRN3X

The only hardware change in the 14th gen stack, but it doesn't do much.

WE'VE HEAVILY DOCUMENTED just how lackluster a launch Intel's 14th generation of processors has been. At its best, it introduced an early look at APO and how Intel, with the help of Windows and Microsoft, can better leverage CPU resources to significantly boost performance across all manner of programs. At its worst, it gave us a 200 MHz overclock on the Intel Core i5-13700K and called it a day.

It is, without a doubt, a bit of a 'meh' generation, hardly worthy of the fanfare of a generational number bump, particularly given we're expecting Intel's 15th series with the launch of the LGA1851 socket later this year. No, this one's been an odd launch, that's for sure. The Intel Core i7-14700K, however, is an intriguing unit, as it's the only chip out of the revised product stack that has had any form of hardware enhancement compared to its 13th-generation counterpart.

You still get those eight performance-cores driving the vast bulk of the performance, but now, instead of a paltry eight efficient-cores, you're graced with a grand total of 12 of them. This boosts the total effective thread count to 28, only four behind Intel's own 14900K. In fact, that's one of the few differences between the two chips. That core i9 has four additional efficient-cores, a slightly higher max turbo frequency (6 GHz to 5.6 GHz), and that's about it. There's not a whole lot more to say than that.

The only other downside is the lack of APO support for any Core i7 model or below, which could have been a saving grace for this generation. At the time of writing, the jury's still out on whether this could be the game-changing AI grace that we hope Intel is going to make it. Early results have been promising, but unless the list of actively supported programs increases, and there's an easier method of installing and activating it (Microsoft's store is still not good), there's a big chance this thing is going to be a wet noodle, rather than the cutting-edge katana we hope it will be. It has the potential to be game-changing, figuratively and literally, but it requires Intel to back it, and given it's only available for Core i9 processors right now—not even for its plucky Core i7 here—that makes the target audience who has access to it rather limited.

That aside, what's the Core i7-14700K like on the day-to-day? Simply put, it's a convenient little chip that packs a hell of a punch. If it wasn't for the clock-speed differences between it and its big i9 brother, there wouldn't be a huge amount to talk about. There's a $160 price difference between the two as well, which is a lot to pay for APO, four efficient-cores, and a 400 MHz bump. Otherwise, it performs as well as you'd expect it to. CineBench numbers are strong, clocking in 2,103 on the single-core test and 32,764 on multi-threaded. Puget bench clocks in a respectable 9,103, and generally performs well in gaming across the board, being within one percent of our Core i9-14900K across almost all titles.

It gets a little trickier when it comes to power and temperatures, however. It's a thirsty and hot little thing, consistently hitting the 100 degree mark under load, even with a 360mm rad on top of it, and in our testing it hit 532.0W under load in 3DMark's Speedway benchmark. That said, we have seen higher loads during real-world benchmarking.

One thing we will point out is that we did have trouble very specifically with our test-bed setup. For some reason, no matter what we did, outside of BIOS on the Asus ROG Maximus Z790 Dark Hero, the 14700K continually locked up Windows. We ran through every possible troubleshooting solution, including wiping the SSD and trying to reinstall Windows (locked up on the install, too), and still no joy. In the end, we had to swap out the Z790 Dark Hero for the MSI MEG Z790 ACE MAX , which did the trick. We're currently working with Intel and Asus to try and figure out why this is issue is occurring, but we'll keep you in the loop.

Testing aside, the 14700K is a funny little chip. In isolation, it produces some stunning performance, as you'd expect. It's hot and power-hungry, but we expect that from Intel's 14th gen at this point, it's just, generationally, a bit dull, which is a shame. Maybe next time, Intel. **–ZAK STOREY**

BENCHMARKS

	Intel Core i7-14700K	Intel Core i9-14900K	Intel Core i5-14600K
CineBench R23 - Single Core (Index)	2,103	**2,194**	1,956
CineBench R23 - Multi Core (Index)	32,764	**35,959**	22,638
Puget Bench - Photoshop (Index)	9,103	**9,799**	7,597
SiSoftSandra - Memory Bandwidth (GB/s)	86.54	**87.44**	82.25
SiSoftSandra - Memory Latency (ns)	71.90	**71.70**	72.50
3DMark - Fire Strike (Index)	42,053	**44,609**	47,710
3DMark - Port Royal (Index)	18,093	**18,107**	18,027
3DMark - Speed Way (Index)	7,278	**7,322**	7,213
Power Draw - Speed Way (Watts)	532.0	545.0	**511.4**
Power Draw - Desktop Idle (Watts)	189.8	190.8	**170.8**
Max Temperature - Speedway (Celsius)	100	100	**84**
Idle Temperature - Desktop (Celsius)	48	**46**	48
Total War: Three Kingdoms (avg fps)	**218**	217	212
Horizon Zero Dawn (avg fps)	210	209	**214**
Cyberpunk 2077 (avg fps)	223	**224**	218

Best scores in bold. Our test bed consists of an Asus ROG Maximus Z790 Dark Hero & MSI MEG Z790 ACE MAX, 32GB of Corsair Dominator Titanium @ 7200, an Nvidia GeForce RTX 4080, Corsair H150i AIO, and an Adata Legend 960 Max PCIe 4.0 SSD. All gaming tests were performed at 1080p, on the highest default preset. XMP is enabled.

VERDICT

8

Intel Core i7-14700K

+ ALL THE GEAR Strong multi-core and single-core performance; Great in-game; Decent pricing; Hardware update on 13th gen.

− NO IDEA Mediocre generational update; Very similar to a Core i9.

$419, www.intel.com

SPECIFICATIONS

Cores (P+E)	8 + 12
Threads	28
L3 Cache (Smart Cache)	33 MB
L2 Cache	28 MB
Unlocked	Yes
Max PCIe Lanes	20
Graphics	Intel UHD Graphics 770
Memory Support (up to)	DDR5 5600 / DDR4 3200
Processor Base Power (W)	125
Maximum Turbo Power (W)	253

Steam Deck OLED

This is what handheld PC gaming is meant to be

IF THE OG STEAM DECK was a beta, the new Steam Deck OLED is its finished form. It's not a sequel, not a beefed up 'pro' version, nor a major redesign. Heck, it's hard to tell the difference between the old new models, save for the new orange accents. But there are differences, and they make, well, all the difference.

For instance, it has better battery life, more storage capacity, and weighs a tad less. Then there's that sharper, brighter screen, upgraded from LCD to OLED and running at 90Hz instead of 60Hz. But there are surprises, too. The touchscreen is way more responsive, the thumbsticks have an improved grip, and the bigger fan purrs more softly than before.

The main result? You won't want to go back to the original Steam Deck. Indeed, that OLED panel captures the promise of the original in ways you might not expect. While upgrading may not be compelling for existing Deck owners, the Steam Deck OLED cements Valve's device as the most accessible handheld gaming PC.

Games positively radiate on the new Deck's slightly larger 7.4-inch screen. The OLED panel can push 1,000 nits of brightness. Take your first steps into Elden Ring's pastoral opening region without properly adjusting your settings, and it's like someone flipped on light mode. With an impressive 110 percent of the DCI-P3 gamut, the colors on the Deck OLED are deliciously vibrant.

Of course, the AMD APU is essentially the same as before, and offers the same four CPU cores and RDNA 2 derived graphics. That's why the resolution is still 1,280 by 800 pixels. The Deck OLED is not a performance upgrade and that's intentionally so. Valve wants to maintain a consistent performance target for developers with the new device.

Depending on how bright you run the new Deck, the cost can be battery life. The Deck OLED bumps the battery up to 50Wh (from 40Wh) for a noticeable improvement in best-case-scenario battery life. With default settings on

the original Deck (60Hz, Wi-Fi on) and default settings on the Deck OLED (90Hz, Wi-Fi on), you can start to see what that new battery is doing for you. The Deck OLED pulled out nearly another 45 to 60 minutes of battery life compared to the original, but enabling HDR and running the screen at full brightness can knock 20 minutes or so off mains-free gaming.

Of course, at 90Hz, the OLED panel runs smoother in games that can achieve that kind of frame rate, and the sub-1ms response means it's sharper and clearer. Of course, not every Deck model gets these enhancements. As part of the Deck OLED refresh, Valve is only keeping the 256GB LCD model around and lowering its price to $399, which is compelling if you can't afford the next step up. A 512GB Deck OLED for $549 now sits in the middle next to a 1TB model for $649.

The most expensive version also comes with a carrying case that hides another, thinner carrying case. Anyway, for those who hadn't already decided on a Deck, the new OLED model is the best version. For anyone with a big Steam library or the drive to tinker with Linux, there's no better handheld available.

When we visited Valve to pick a Deck OLED up, almost every developer I met described the Deck OLED as the device they wanted to make from the beginning. It's not hard to see why. That it's available at a price that still slips in under the competition despite its improvements sweetens the deal, and ensures the Steam Deck OLED retains the king of handheld gaming PCs crown. **–TYLER COLP**

VERDICT
9
KICK ASS!

Steam Deck OLED

+ SPECTACULAR SEQUEL Stunning OLED screen; Several detail improvements.

– SECOND RATE Not actually any faster; Battery life better, but still fairly limited.

$549, www.steampowered.com

SPECIFICATIONS

APU	AMD Sephiroth
Coresa	4
Threads	8
GPU	RDNA 2
Compute units	8
Memory	16GB LPDDR5-6400
Storage	512GB NVMe SSD
Screen	7.4-inch OLED
Peak luminance	1,000cd/m² (HDR) 600cd/m² (SDR)
Native resolution	1280 x 800
Refresh rate	90Hz
Network	Tri-band WiFi 6E, BT 5.3
Battery	50Whr

It all looks very familiar, but the Deck OLED has been thoroughly revised.

RCA Evolution Premium M27PG135F

Blast from the past

DO NOT ADJUST YOUR SETS; your eyes do not deceive you. This is a high-refresh gaming monitor from RCA, a company that pioneered early radio and TV tech, and is now essentially a piece of branding IP licensed by several companies.

In this case, it's DTI Services, a company that primarily makes TVs, but is now branching out into PC gaming monitors. Enter the new RCA Evolution Premium M27PG135F, a 27-inch 1440p 240Hz model, and our first taste of the brand's new foray into gaming hardware.

Out of the box, it's pretty familiar. RCA isn't trying to reinvent the gaming monitor with the M27PG135F. The build quality is a little hollow and plastic, but that's typical for this class of display. You get a reasonably sturdy stand with a metal base and a full range of adjustments, including pivot into portrait, while the slim-ish bezels on three sides are standard aesthetic fare, as is the light sprinkling of RGB lighting on the rear.

Connectivity is well covered, too, with DisplayPort, two HDMI sockets, and even USB-C with 90W of power delivery. The inclusion of a KVM switch is a nice addition, and contributes to a pretty strong all-round feature set.

Specs-wise, there are no surprises. This is a 27-inch 1440p model with 240Hz refresh capability, and 1ms claimed GtG response. There is basic HDR support, but with a peak brightness of 300 nits, it doesn't qualify for explicit DisplayHDR 400 certification.

But what of the image quality? It's decent, but not remarkable. The color calibration is on point, and the subjective brightness is beyond expectations, given that 300 nit rating. Predictably, this monitor doesn't give a true HDR experience. There's no local dimming—this is a conventional IPS panel with limited contrast performance.

So while HDR video doesn't look bad, it's debatable if it looks better than in SDR mode. It's more complicated with gaming. Running, say, *Cyberpunk 2077* with HDR switched on, brighter highlights have more pop than in SDR mode.

More important is speed, where it puts on a stronger showing. There are three levels of pixel accelerating overdrive in the OSD menu, and even the fastest option only introduces a tiny level of overshoot.

You might think that's good. However, a more aggressive max setting can be desirable. If you're playing, for instance, *Counter-Strike 2*, that responds well to aggressive overdrive. It tends to look crisper and sharper, the overshoot and inverse ghosting not noticeable.

RCA could have been less conservative with the overdrive options. On the other hand, we doubt you'll have any problems with latency. All 240Hz are accounted for, making for buttery-smooth frame rates and essentially zero subjective latency. All in all, it gets pretty much nothing badly wrong and a lot right.

The only slight snag is the question of why you'd buy this monitor over the more obvious brands. It's a decent performer, but so is much of the competition. It's reasonably priced, but then so is... well, you get the idea. This is a solid first effort from RCA, but the brand probably needs to do a little more to earn conquest sales from the better established competition.

–JEREMY LAIRD

VERDICT
8

RCA Evolution Premium M27PG135F

⊞ **CLASSIC** Speedy IPS panel; Decent all-round feature set.

⊟ **CLUNKY** Slightly limited brightness; Only basic HDR support.

$369, www.rcamonitors.com

SPECIFICATIONS

Screen size	27-inch
Resolution	2,560 x 1,440
Brightness	300 nits
Color coverage	95% DCI-P3
Response time	1ms
Refresh rate	240Hz
HDR	HDR10
Features	IPS panel, adaptive sync, 1x DisplayPort, 2x HDMI, 1x USB-C with 90W PD, KVM switch

© UNSPLASH

An old brand from yesteryear has learnt some new tricks.

Hyte Y70 Touch

Almost perfect, but not quite

IT'S FAIR TO say that Hyte's Y70 Touch has impressed us—it's a chassis unlike any we've seen before. That angular design, touchscreen display, and rich feature set has made it one of the nicest cases we've had the pleasure to work in. There's a lot to love here—a lot of sweet little touches that make the Y70 Touch a thing of beauty. However, it's not without fault. Throughout our time building with it, there were hiccups; moments where we questioned design choices, decisions that baffled the mind, and slight build inconsistencies that lack the refinement of a more mainstream chassis builder.

Let's start with the good. It's a full-sized ATX case, capable of holding E-ATX boards in its girthy housing, complete with some serious cooling potential. In total, you can mount ten fans in here at any one time (not including push-pull): three in the floor, three in the roof, three in the side, and one in the rear. There's plenty of space for radiators and liquid-cooled solutions, and the overall chill factor is demonstrable.

On top of that, you get an impeccable dual-chamber design, super-slick rubber grommets, and a vertical GPU mount as standard. For build features, there are removable radiator brackets, support for 2.5 and 3.5-inch drives (albeit only two of either), and plenty of cable management room to play around with. If you do decide to go all the way, you can strap those pesky things down, or shove them all in the back of the case.

Then there's the display: 1,100 x 3,840 pixels of pure 14.1" gloriousness, complete with a 60Hz refresh rate and 10-point multi touch support. It's all powered via USB and DisplayPort. It looks crisp, clean, and acts exactly as you'd expect. Hook it up to your GPU, and watch as Windows recognizes it as exactly what it is: an additional display. Sling Slack into it, or Reddit or Discord, and it works like a treat. Of course, you'll get the most mileage out of it using Hyte's Nexus app (a neat little program that allows you to bake all manner of widgets and system specs into it), but all in all, it's an incredibly nice addition, and it even gives the chassis an almost reasonable 'value' angle.

As for negatives, the first is the price, clocking in at $359.99. Don't get us wrong, this is a stunning case, and that extra display adds an element of pizzazz. Couple that with the features we mentioned, and the PCIe 4.0 riser cable ($80 on its own), and it suddenly becomes a lot more palatable.

However, you can't buy this specific variant without that display. The Y60 exists, but it's slimmer, doesn't have quite the cooling capacity with its vertically mounted GPUs, and just looks a bit chunkier around the 'bezels'.

Then there's the GPU problem. Specifically, you can only mount a GPU on that vertical mount. There's no option to enable you to install it horizontally, as all the PCIe slots are half-height, meaning similarly on the price front, you're locked into getting that PCIe riser. That works fine, but it's the lack of choice that really frustrates us.

It's not all rose petals and dew drops. The bottom fan bracket technically supports three 120mm fans, but in our case, we just couldn't get three to mount securely down there, so they're loose. On top of that, the paint finish, although a fantastic powder-coat, has that orange-peel texture of a car that hasn't been polished correctly, which again, given the price, we just don't expect.

Still, all that aside, and ignoring the lack of choice, the Hyte Y70 Touch is an impressive chassis. It's redefined what a case can be, particularly in this age of boring fish-bowl designs or black towers, and for that we have to give it credit. Is it perfect? No, but it doesn't have to be. No case ever will be, and what it does right, it absolutely nails. If you pick one of these up, you'll be more than happy with it sitting proud on your desk. **–ZAK STOREY**

VERDICT

9

Hyte Y70 Touch

➕ **UPPER ECHELONS** Outlandish design; Incredible display; Intuitive build features; Vast cooling options, including PCIe riser; Glorious aesthetic.

➖ **LOWER PITS** No budget alternatives; Limits you to vertical GPUs; Can't 'mount' three bottom fans; Paint-work needs a touch-up.

$359.99, www.hyte.com

SPECIFICATIONS

Motherboard Support	ITX, Micro-ATX, ATX, E-ATX 2.5-inch / 3.5-inch Support: 2x 2.5/3.5-inch
Max Radiator Support	360mm Side, 140mm Rear, 360mm Roof
Fan Support	1x 120/140mm Rear, 3x 120 / 2x 140 Side, 3x 120 / 2x 140 Roof, 3x 120 / 2x 140 Floor
Dimensions	18.5 x 12.59 x 18.5 - inches
Graphics Card Clearance	16.6 x 4.1 - inches
CPU Tower Clearance	7.0 - inches
Warranty	Three Years

CPU Core #1
Thread #1
11%
CPU Core #1
Thread #1
11%
Used Space
83%
Memory
28%
GPU Core 42%
GPU Core 30°C
GPU Core 210MHz
54°
Clear
H: 71.8° L: 50°
Los Angeles, California
10:37pm 8:44am
HYTE
A ATX
M Micro-ATX
Mini-ITX

There are colorways galore with this thing.

MEG
ACE
MAX
LGA-1700

Great compared to the competition, but not so much against its cheaper siblings.

MSI MEG Z790 ACE MAX

The budget high-end Z790 option… sort of

IT'S NO MYSTERY that here at *Maximum PC*, we aren't exactly fans of super-expensive motherboards. In a world where bang for your buck is king, $700 super mobos that typically shouldn't act more than a fancy interconnect aren't what we'd consider 'good investments' long term. After all, these are exactly that—interconnects that are going to house the latest and greatest components. They are without a doubt luxury items, and the MSI MEG Z790 ACE MAX is exactly that—filled to the brim with the latest gadgets and gizmos built to woo that hard-earned cash from your wallet.

As a unit, this thing is a monster. At a retail price of just $700 (on par with the Asus ROG Maximus Z790 Dark Hero), there are very few boards out there that pip this thing on a price front, and it shows. This is absolutely brimming to the hilt with added features and extras you're not likely to find for $300 (not in the last couple of generations of chipsets, anyway). There's a 24+1+2 power phase design, dual 2.5Gb lan ports, WiFi 7 as standard, impressive thermal cooling solutions, aluminum backplates, a plethora of M.2 ports (including twin PCIe 5.0 slots + a PCIe 5.0 M.2 slot), and support for DDR5 up to 7,800 MT/s.

Take a look at the rear I/O, and you'll get a sense of scale for this thing. Seven USB 3.2 Type-A ports, one Type-C port, two Thunderbolt 4 Type C, WiFi 7, two Mini DisplayPort, 5.1 Audio + S/PDIF out, flash BIOS, Clear CMOS, and a 'Smart' button (you can configure this in BIOS to either reset the PC, turn on/off Mystic Light, go into Safe Boot, or run all fans at 100 percent/Default speeds).

What's not so smart, however, is the aesthetic. Sure, the triangles are nice, but we're still greeted with MSI's bizarre dragon motif, along with some rather basic RGB lighting across the board. The brushed aluminum is a nice touch, but again, doesn't feel quite as premium as some of its sandblasted competition, or mirrored opposition. That's not the be-all and end-all, of course, but at $699, it does hold some significant weight here in the field of glorified interconnects.

Where the Ace Max wins out is on the performance front. It wipes the floor with the Asus ROG Maximus Z790 Dark Hero, its closest competitor, both in our benchmarks and from a pure interconnect standpoint. It has more on-board ports rear I/O, better power phase design, and a better board layout. It runs rings around the opposition, making it a far better pick for those professionals looking to max out what they connect to their motherboard. The twin 2.5Gb ports alone are worth pointing out.

Here's the thing, though. We've also managed to get MSI's own Z790i EDGE WiFi board through the benchmark suite, a board that's well under half the cost of the Z790 ACE MAX. Perhaps unsurprisingly from a purely numbers perspective, with the same hardware, we're looking at similar performance across the board. The ACE MAX continues to pip right past it, but it also costs twice the cash. Once again, price isn't everything.

MSI's Z790 ACE MAX is an intriguing product. On the one hand, compared to motherboards at this price point, it's an incredibly impressive proposition, one that provides far better connectivity and performance than its counterparts. On the other, even compared to its own smaller, cheaper siblings, performance remains similar. If you absolutely need the connectivity, and like the idea of a board that'll future-proof you for some time to come (albeit there's no guarantee 15th Gen Intel will fit in an LGA1700 socket), then it's a fine pick. If you just want to save a bit of cash, you're far better off splurging for something a little (ok, a lot) cheaper instead. **–ZAK STOREY**

BENCHMARKS

	MSI MEG Z790 Ace Max	Asus ROG Maximus Z790 Dark Hero	MSI Z790i Edge WiFi
GeekBench Single / Multi (index)	**2,797** / **17,646**	2,793 / 17,560	2,791 / 17,524
CineBench R23 Single / Multi (Index)	**1,960** / 22,890	1,948 / **23,109**	1,953 / 22,790
CrystalDisk Mark Sequential Q32T1 (MB/s)	**7,133** / 5,987	7,131 / **6,113**	7,115 / 6,006
CrystalDisk Mark Random 4K Q1T1 (MB/s)	79.93 / 293.29	**89.38** / **329.90**	79.46 / 312.53
SiSoftSandra Memory Bandwidth (GB/s)	76.86	**83.48**	77.85
SiSoftSandra Memory Latency (ns)	73.6	**72.2**	74.5
CPU / GPU Temperature Desktop @ Load (C)	90 / 64	83 / **63**	**82** / **63**
CPU & GPU Power Draw @ Load (watts)	483.03	476.66	**469.45**
3DMark: Fire Strike (Index)	47,162	**47,196**	47,032
Horizon Zero Dawn	**227**	220	219
Total War: Warhammer III	**218**	209	207
Cyberpunk 2077 RTX	**125**	**125**	**125**

Best scores in bold. Our test bed consists of an Intel Core i5-14600K, 32GB of Corsair Dominator Titanium @ 7,200, an Nvidia GeForce RTX 4080, Corsair H150i AIO, and an Adata Legend 960 Max PCIe 4.0 SSD. All gaming tests were performed at 1080p, on the highest preset. XMP is enabled.

VERDICT

8

MSI MEG Z790 ACE MAX

+ **ACE IN THE HOLE** Expansive connectivity; Better power solution than the competition; Impressive overall material use; Strong performance.

− **ACE VENTURA** Dragon logo sorely needs an update; Far cheaper options provide similar performance.

$699, www.msi.com

SPECIFICATIONS

Chipset/Socket	Z790 / LGA1700
Max Memory Support	192GB DDR5 @ 7,800+
Storage Support	5x M.2 Slots (1x 5.0x4 and 4x 4.0x4), 6x SATA 6Gb/s
PCIe Support	2x PCIe 5.0 x16 (x8x8), 1x PCIe 4.0 x16
Form Factor	E-ATX, 12 x 10.9-inches
Rear I/O	2x Thunderbolt 4 Type C, 8x USB 10Gbps (7x Type A, 1x Type C), 2x Mini DisplayPort (Thunderbolt), 1x WiFi 7 Antenna, 2x Intel 2.5Gb Port, 5.1 Audio + optical S/PDIF Out, BIOS FlashBack, Clear CMOS, Smart Switch

A great touchscreen monitor, if touch is what you need.

Asus ZenScreen Ink MB14AHD

A portable touchscreen for the traveling creative mind

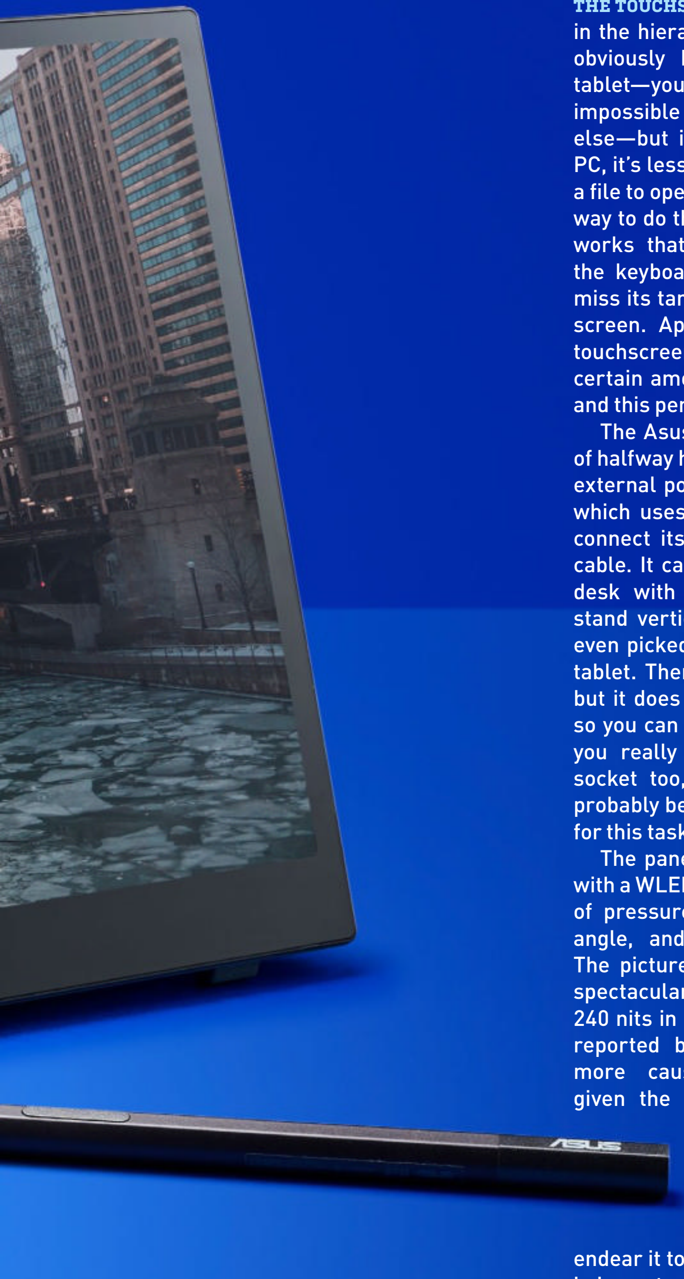

THE TOUCHSCREEN has a precarious place in the hierarchy of PC input devices. It's obviously brilliant on a cellphone or tablet—you can see this in the way it's impossible to buy one that uses anything else—but in the world of the Windows PC, it's less clear. Sure, being able to tap a file to open it sounds like a much easier way to do things, but in practise it rarely works that way, having to reach over the keyboard, only to have your finger miss its target and leave a smear on the screen. Apple hasn't bothered making touchscreen MacBooks despite having a certain amount of expertise in the area, and this perhaps tells us something.

The Asus Zenscreen Ink offers a sort of halfway house, being a touch-sensitive external portable monitor with a stylus, which uses the cleverness of USB-C to connect itself to your PC with a single cable. It can be set horizontally on your desk with its built-in stand, made to stand vertically with a bit less tilt, and even picked up and used like a graphics tablet. There's no wireless connectivity, but it does have a pair of tiny speakers, so you can use it for watching movies, if you really must—there's a headphone socket too, although most people will probably be using Bluetooth headphones for this task.

The panel is a 14in 1080p IPS model, with a WLED backlight. It has 4,096 levels of pressure sensitivity, a 170° viewing angle, and 10-point touch sensitivity. The picture is even and sharp, but not spectacularly bright, peaking at around 240 nits in our tests. The color accuracy reported by our colorimeter provides more cause for concern, however, given the sort of people likely to be interested in the ZenScreen Ink. It displays 62 percent of sRGB, 47 percent of AdobeRGB, and 47 percent of P3, which are not going to endear it to those for who color accuracy is important. Of course, it's unlikely that anyone would use the ZenScreen as their primary monitor, instead teaming it with something larger and better calibrated, but in a device aimed at the creative industries, it would be nice to see better color response.

The ZenScreen Ink comes with a stylus, which is fully supported by Microsoft Pen Protocol 2.0 over USB-C in Windows. It's a stubby black pencil-like model with a button at the exact spot your finger would naturally fall on. Its only drawback is that it requires a battery, like we're back in 1997 or something. Worse, it's an AAAA cell, a kind we weren't sure even existed. You can pick them up for a few bucks, but a rechargeable stylus just seems like a better solution, though we appreciate that it's quicker to just swap a battery over than wait for it to charge.

The ZenScreen Ink also comes with a nicely made sleeve, its gray fabric doing an excellent job of protecting your touchscreen. Alongside the built-in stand, there's also a screw thread for connecting it to a camera tripod, something that may draw those in photo studio environments to it, though the OLED model we reviewed previously would be a better fit there, given its superior color reproduction.

Otherwise, while it's nice to have touch-sensitivity in an external screen, it's debatable how much it will get used. If you're already using a touchscreen, then it can prevent that embarrassing moment when you try to scroll on a screen that doesn't support it, and the stylus is good for drawing on a mirrored screen as if it were a graphics tablet. But whether it's worth a premium over a non-touch model remains questionable. **–IAN EVENDEN**

VERDICT
7

Asus ZenScreen Ink MB14AHD

➕ **INKY** Sharp display; Touch sensitive; Useful stylus.

➖ **STINKY** Narrow color response; AAAA battery.

$299, www.asus.com

SPECIFICATIONS

Screen	14in 1080p IPS WLED, 60Hz
Max Brightness	220 nits
Colour depth	8-bit
Response time	5ms
Ports	1x Micro-HDMI, 2x USB-C (DisplayPort 1.2 Alt Mode), tripod socket
Pressure sensitivity	4096 levels
Touch sensitivity	10 points
Dimensions	12.76 x 7.92 x 0.49 inches
Weight	1.91 pounds

© UNSPLASH

SUBSCRIBE TODAY

Find your next great read

On iOS & Android!

Z790 MOBOS REVIEWED
WINDOWS COPILOT FULL GUIDE
HOW TO BAN ADS
MAXIMUM PC
BUILD A TOUCH SCREEN GAMING PC
Core i5-14600K + RTX 4080
4K Gaming Powerhouse
14th Gen i5 vs i9 Compared
LIVESTREAM LIKE A PRO
WIN 10: STICK OR SWITCH?
FIND EVERY HIDDEN FEATURE IN macOS SONOMA
Mac|Life
THE NEW M3 MACBOOK PRO
CITIES: SKYLINES II REVIEWED
PC GAMER
ROMERO RETURNS
DOOM
GAME OF THE YEAR
ALAN WAKE 2
NIGHTINGALE

Subscribe today and get instant access on your iPad, iPhone or Android device.

www.magazinesdirect.com/B3XMPC

32GB XPG Lancer Blade RGB DDR5

Does memory speed matter?

MEMORY FREQUENCY has long been the realm of marketing gurus and PR specialists, those whose task it is to upsell RAM and keep that monetary train flowing. After all, in every other component, frequency has a significant impact on performance and progress. By that definition, it's surely logical that memory follows the same route?

Not quite. For a long time, Intel's processors benefited little from memory speed. It was only AMD with its Ryzen processors, which tied its own PCB interconnect to memory speed directly, that led to frequency being worthwhile for CPU performance. That's not to say memory speed wasn't warranted, just that capacity and real-world latency were king. If you were rendering massive video files or dabbling in complex database tasks, frequency did have an effect.

Today, we're reviewing something different: Adata's XPG line of Lancer Blade RGB DDR5 memory, comparing it directly with Corsair's Dominator Titanium memory SKU. Both are 32GB at 2x16GB, DDR5, and have been tested using Intel's Core i5-14600K processor on an Asus ROG Maximus Z790 Dark Hero motherboard. With capacities and test-beds the same, does the 6000 MT/s kit from Adata pip Corsair's 7200 kit?

Real-world latency is the metric we need to keep an eye on when it comes to RAM performance. Frequency only tells one story, and the reality is that it's a combination of frequency and CAS latency that really dictates frequency. You can calculate this with any memory kit, as long as you have those two numbers. Here's the formula: 2,000 divided by memory frequency, multiply that by your CAS latency, and you'll end up with your calculated real-world latency in nanoseconds. In our case, the XPG Lancer Blade (with 6,000 MT/s and a CAS latency of 30) clocks in at 10 ns, our Corsair kit (7,200 MT/s and a CAS latency of 34) comes in slightly faster at 9.44 ns.

How does that affect performance? In our 10GB WINRAR Archive Time, the difference was one second (with the Titaniums coming out on top), in *Total War Warhammer III* @ 1080p, average frame rates were identical, with the XPG Lancer Blade kit, losing the minimum fps test by 1 fps (147 vs 148).

Flip over to synthetic tests, and things get a little gray around the edge. SiSoft Sandra's software suite clocks wins across all fronts for the Dominator Titanium kit, as does PCMark10, and PugetBench's Photoshop testing run scores it 600 points higher as well.

Here's the trick. The two kits are different when it comes to price, and the gigabyte per $ metric is huge, with the XPG Lancer clocking in 0.3GB per $, versus the Dominator Titanium's 0.15GB per $. That's double the cost per gigabyte. Of course, you get better performance in those rendering and editing tasks, plus access to iCUE, among other things.

So what's the answer? If you're on a tight budget, unless you're doing very specific video editing or rendering tasks, for day-to-day usage, your best bet is to buy a reasonably priced kit, with a focus on capacity. If you're running an AMD 7 series, we'd recommend a 6,000 MT/s kit at most. In this case, Adata's XPG Lancer Blade RGB is a true value king, with an alluring design, strong performance, and a sleek enough profile to slide right into any build. **–ZAK STOREY**

BENCHMARKS

	32GB (2x16GB) XPG Lancer Blade RGB DDR5 @ 6000 MT/s	32GB (2x16GB) Corsair Dominator Titanium DDR5 @ 7200 MT/s
SiSoftSandra Overall Memory Score (Index/kPT)	2.27	**2.49**
SiSoftSandra Memory Latency (ns)	80	**73**
SiSoftSandra Memory Bandwidth (GB/s)	69	**80**
PCMark10 - Express (Index)	6,624	**7,162**
10GB WinRAR Archive Time (Seconds)	116	**115**
Puget Bench - Adobe Photoshop (Index)	7,550	**8,150**
Total War: Warhammer III (avg/min fps)	**191.3** / 147	**191.3 / 148**
Max Temp Under Load (Celsius)	51.0	**43.5**
Real World Latency (ns)	10.00	**9.44**
Gigabyte per $ (Index)	**0.30**	0.15

Best scores in bold. Our test bed consists of an Intel Core i5-14600K, Asus Z790 Dark Hero motherboard, an Nvidia GeForce RTX 4080, Corsair H150i AIO, and an Adata Legend 960 Max PCIe 4.0 SSD. All gaming tests were performed at 1080p, on the highest preset. XMP is enabled.

VERDICT
9

32GB XPG Lancer Blade RGB DDR5

⊞ **MASTER SWORD** Impressive value option; Subtle RGB lighting; Good all-round performance; Available in black and white.

⊟ **BROKEN HILT** Not the absolute best in editing tasks; Software could be better.

$104.99, www.xpg.com

SPECIFICATIONS

Capacity	32GB (2x16GB)
Channels	Dual
DDR Standard	DDR5
Frequency	6,000 MT/s
CAS Latency	34
Operating Voltage	1.35-1.4V
Dimensions	133.35 x 40 x 7.86mm
Operating Temperature	0°C - 85°C
Warranty	Limited Lifetime

Sony H9 Wireless Noise Cancelling Gaming Headset

That's legitimately its name

WE'RE COMING IN all guns blazing with this one. It's not good, and honestly, that's a darn shame, particularly given Sony's awesome heritage when it comes to both top-tier audio quality and noise cancelling.

The task was simple for Sony: take the people down the lab who worked on the M4 series of noise-cancelling headphones. You give them a bit of RnD budget, say, "Build us a kick-ass gaming headset for PC and PS5", and leave them to it. That's all you needed to do. Instead, we've got this mediocre thing.

Is it that bad, then? Well, let's reel off the top line. It's a $250 headset, designed specifically for PS5 and PC usage. It's got a 32-hour battery life, is wireless, big, and bulky, with some plush foam ear-cups, and on paper, some decent stats. 40mm drivers, closed-dynamic design, active noise cancelling, Bluetooth 5.0 support, 10-minute quick-charge, and an embedded mic, as well as a 5 Hz to 20,000 Hz frequency response. On paper, it's all there.

It looks the part, too. If you already own a PlayStation 5, you'll be used to the DualSense aesthetic this thing eschews: the sandblasted white plastic finish, the crisp, clean branding, super-stylish black plastic ear-cup holders, and cleverly built-in noise cancelling mics throughout. If you're looking to complete your PS5 gaming setup, this thing certainly fits the part.

Connectivity-wise, you've got that Bluetooth 5.0 support we already mentioned, but it also comes with 2.4 GHz wireless as well, courtesy of a dongle. Plug it into your PS5, chuck the switch on it to 'PS5', and you're good to go. If you need it on PC, flip the toggle the other way to 'PC', and voila, you're capable of connecting to your PC.

The downfall, perhaps unsurprisingly, comes from the audio quality, or lack thereof, and its a real shame. Overall, the H9 just lacks the crisp clarity you'd expect from a headset at this price point. Even for a wireless unit, something like a Virtuoso, or Logitech's Pro X 2 Lightspeed, just runs rings around it. Admittedly, they don't have the noise-cancelling clout, but that's not why you buy a headset; you buy a headset for the audio quality. In the H9's case, it lacks the top trebles, the mids are fairly muted, and the bass, although prominent, lacks the punch we'd expect from a unit at this price point.

You can configure your way out of some of those issues using Sony's legitimately well-developed software suite, but it requires fairly extensive EQ tuning to come close to that, and it still lacks the bass oomph we'd expect, which is surprising given the range these drivers can pump at.

As for noise cancelling, it's good at eliminating low-level humming and muting sound in that range, but that's it—it's not quite as good as Sony's own M4s, or at dealing with audible voices, or anything of that ilk. What's more impressive is actually the ambient sound mode, which turns the H9's into more of an open-backed esque style of design, albeit without the roomy feel you get from those kinds of cans. But still, in a gaming environment (which these are specifically designed for, given the fixed mic), it's just a bit of a gimmick.

Really, you're going to be buying this headset if you want the 'Sony' badge clout, and for it to match your DualSense controller. There's nothing here for an established audiophile or anyone who even mildly enjoys good-value, good-quality audio, which is such a shame. Is it an enjoyable experience? Yes. But is it good for the price? No. **–ZAK STOREY**

VERDICT **7**

Sony H9 Wireless Noise Cancelling Gaming Headset

⊞ **HIFI APPROVED** Beautiful design elements; Good software suite; Comfortable.

⊟ **BUY ON SALE** Form over function; Poor audio quality; Lacks audio range; Noise cancelling is gimmicky and flawed.

$248, www.sony.com

SPECIFICATIONS

Driver	40mm Dynamic
Magnet	Neodymium
Frequency Response	5 Hz - 20,000 Hz
Battery Life	32 hours
Connection Type	Wireless 2.4 GHz, Bluetooth 5.0
Platforms	PC, PS5
Mic	Integrated boom arm
Weight	11.64 ounces

WD Black SN770M PCIe 4.0 M.2 SSD

WD_BLACK
SN770M
WDBDNH0010BBK
1TB
S/N: 23443P800936
RN: MVBSN1
3.3VDC 1.8A
Made in Malaysia
05NOV2023
PSID:

Tiny package, big speeds, massive temperatures

HONESTLY, that strapline is perhaps the best summation this journalist has ever written for a product. Western Digital's Black SN770M is designed with one thing in mind: to massively expand your handheld device's storage. Yep, we're talking Steam Decks, Asus ROG Allies, and the numerous new devices launching at CES and Computex this year. The handheld PC has arrived, this time with more graphics and minuscule amounts of internal storage. We bet SSD manufacturers are counting their lucky stars regarding the latter.

Here we have Western Digital's WD Black SN770M PCIe 4.0 NVMe M.2 SSD, and it is seriously tiny. It's the M.2 2230 form factor, and comes in at 29.97 x 22.1 x 2.29mm. It's also available in capacities ranging from 500GB all the way up to 2TB, and to be honest is impressively priced too, ranging from $80 for the 500GB variant, all the way up to $220 for the 2TB model. Ours comes in at $110, giving you 9.17 GB per $ spent, which for a PCIe 4.0 drive of this caliber (and form factor, for that matter) is incredibly well positioned in the market.

As for performance, it holds its own. WD has it listed at sequential reads of 5,150 MB/s and 4,900 MB/s. In our testing CrystalDiskMark 8 topped that, clocking 5,222 and 4,968 respectively. AS SSD was a little less forgiving, clocking 4,550 and 3,818 on sequential read and write, but nonetheless is well within parameters for a solid PCIe 4.0 drive. Access times also performed at the top tier level, clocking in a svelte 0.017 ms for both read and write, even pipping the Gigabyte Aorus Gen5 12000 on the Write access. We suspect that this is down to the design of the drive itself. Given the compact nature of this thing, having both the memory controller and a single NAND chip sat so close to one another should minimize read and write request times considerably, as the drive doesn't have to go searching as far as some of the other drives in our test-bed, giving it an advantage over longer SSDs.

However, this thing is scorching, again no doubt down to the limitations of that size element. Let's be clear, PCIe 5.0 drives are considered hot, so hot that JEDEC recommends each has an active heatsink included by design. You'll likely see either massive heat pipes on them or some form of passive cooling solution (see this issue's Centerfold for a good example). In our case, we're testing all of our SSDs in an Asus ROG Maximus Dark Hero motherboard, with a chunky aluminum heatsink, plus thermal padding. So then, that silly super-hot PCIe 5.0 drive sat under the motherboard heatsink? Maximum temperature, 78 C. It's hot compared to the Samsung PM981 (PCIe 3.0's) maximum temperature of 66 C. But the WD Black SN770M PCIe 4.0 drive? 89 C. That is phenomenal, and just in this testing environment, not large file transfers or access procedures.

Then there's installation. Have you seen how fiddly it is to install an SSD like this into a Steam Deck? It's easier to splurge for a ROG Ally, but even so, it's worth bearing in mind.

WD's SN770M is an SSD of many faces. It's affordable, performs well, and is an impressive feat of engineering. On the other hand, it's hot as hell, and actually installing it in a Steam Deck is a bonafide spec-ops mission. It's unique, and we kind of love it for it. **–ZAK STOREY**

VERDICT
8

WD Black SN770M PCIe 4.0 M.2 SSD

+ **EXTERMINATUS** Solid pricing; Impressive performance; Impeccable engineering.

− **ROLLING 1 TO HIT** Incredibly hot; Challenging to install.

$110, www.westerndigital.com

BENCHMARKS

	1TB WD Black SN770M PCIe 4.0 M.2 SSD	2TB Gigabyte Aorus Gen5 12000 PCIe 5.0 M.2 SSD	1TB Samsung PM981 OEM PCIe 3.0 M.2 SSD
AS SSD Sequential Read / Write (MB/s)	4,550 / 3,818	**8,970 / 9,948**	3,020 / 2,444
AS SSD Random 4K Read / Write (MB/s)	67.31 / 264.00	**86.65 / 289.12**	51.88 / 193.69
AS SSD Access Time (ms)	**0.017 / 0.017**	**0.017** / 0.037	0.076 / 0.019
CrystalDiskMark Sequential QD32 Read / Write (MB/s)	5,222 / 4,968	**12,353 / 11,598**	3,465 / 2,390
CrystalDiskMark Random 4KQ1 Read / Write (MB/s)	**90** / 294	89 / **310**	65 / 201
Max Temp Under Load (C)	89	78	**66**
Gigabyte per $ (GB)	**9.17**	7.69	7.75
Sequential Read MB/s per $ (MB/s)	**47.91**	47.51	26.86

Best scores in bold. Our test bed consists of an Intel Core i9-14900K, 32GB of Corsair Dominator Titanium @ 7200, an Nvidia GeForce RTX 4080, Corsair H150i AIO, and an Asus Z790 Dark Hero. Max Temp recorded via HWMonitor during benchmarking process.

SPECIFICATIONS

Variants	Bare
Form Factor	M.2 2230
Interface / Protocol	PCIe 4.0 / NVMe 1.4
Flash Memory	TLC 3D NAND
Sequential Read	5,150 MB/s
Sequential Write	4,900 MB/s
Random Read	740K IOPS
Random Write	800K IOPS
Endurance (TBW)	600
Part Number	WDBDNH0010BBK-WRSN
Warranty	5 Years Limited Warranty

It's good when it gets going, but boy does *Prince of Persia: The Lost Crown* take its time.

ACTION ADVENTURE GAME

Prince of Persia: The Lost Crown

An often unremarkable attempt at a metroidvania

A SLOW BURN TO A FAULT, *Prince of Persia: The Lost Crown*'s early hours are mired by overly simple and unremarkable platforming, not to mention a disjointed narrative that kind of kills your investment in the story. Based on the first 15 hours alone, we'd be very sour. Thankfully, the game's second half does a much better job, but by the time the credits roll, you can't help wishing that the entire game had been this way.

At its core, *The Lost Crown* essentially takes a tilt at metroidvania, a portmanteau of the names of the videogame series *Metroid* and *Castlevania*, and a sub-genre of action-adventure games focused on guided non-linearity and utility-gated exploration. It's also the *Prince of Persia* franchise's first appearance in a long 14 years. Not only has it swapped out the 3D perspective of its last instalment to return to the franchise's 2D platforming roots, it's also forgone the eponymous prince character in favor of new protagonist, Sargon. This guy isn't royalty, but rather a member of the Immortals, a group tasked with protecting the prince Ghassan and keeping Persia safe.

At least, I'm pretty sure that's the premise of the story. *The Lost Crown*'s narrative is incredibly underbaked. It feels like pieces of the worldbuilding and overarching story are missing. Outside of the Immortals' members, it's hard to follow the characters, their motives, or the routinely baffling and poorly presaged dialogue. Moments that were surely intended to be grand reveals consequently fall flat and as an alternative to constant confusion, the temptation is to opt out of paying attention to the story entirely.

It's a shame, because a stronger narrative would help get through the slog of its early hours. Initially, you have simple jump, slide, and ground-dash moves at your disposal. The training wheels stay on for far too long, and it feels like hours before new abilities eventually come in the form of Time Powers, which allow you to airdash, create copies of Sargon to teleport back to, grab objects like bombs, flick between dimensions to make platforms pop in and out of existence, grapple, and double jump.

As you collect ever more Time Powers, *The Lost Crown*'s platforming creativity is finally unleashed. The real kicker is double jump, for which you have to wait a yawning 15 hours, and at last allows the game to come to life. The difficulty spike in puzzles and platforming with double jump in your toolkit is truly whiplash-inducing.

You go from flying through each area with ease to suddenly hitting wall after wall, whether that be from tight and tricky execution or a puzzle that leaves you head scratching over the powers needed to complete it. This is *The Lost Crown* at

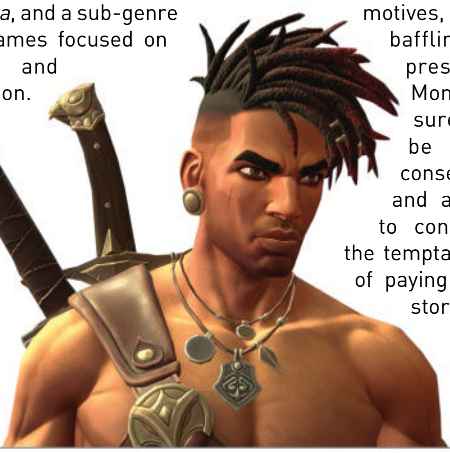

© UBISOFT

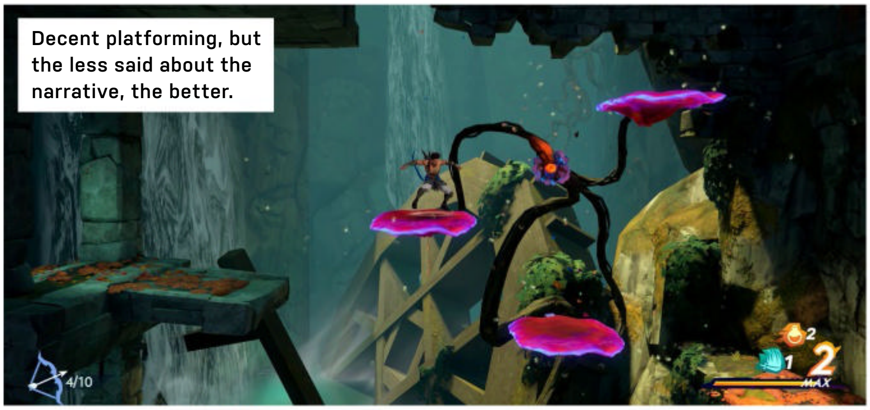

Decent platforming, but the less said about the narrative, the better.

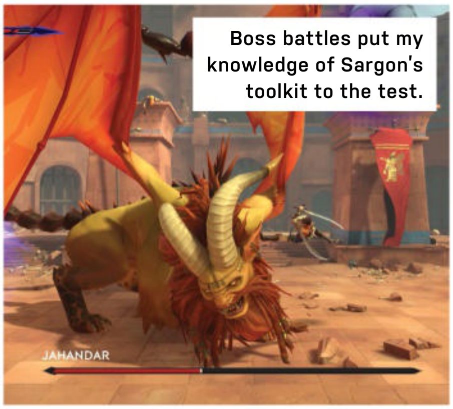

Boss battles put my knowledge of Sargon's toolkit to the test.

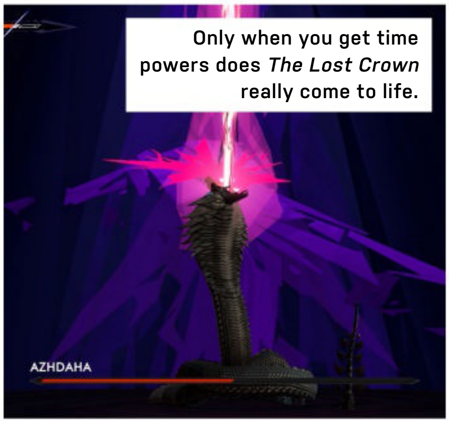

Only when you get time powers does *The Lost Crown* really come to life.

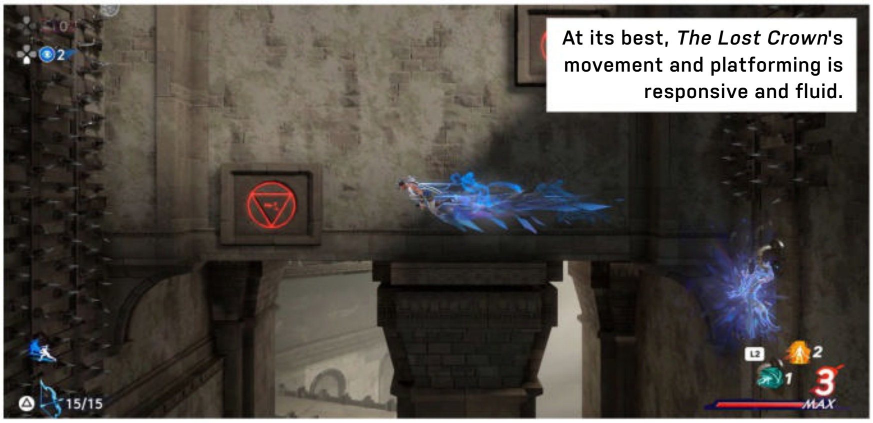

At its best, *The Lost Crown*'s movement and platforming is responsive and fluid.

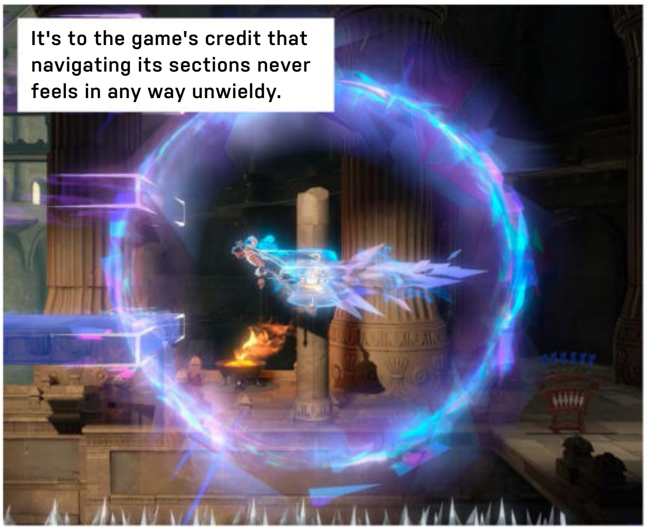

It's to the game's credit that navigating its sections never feels in any way unwieldy.

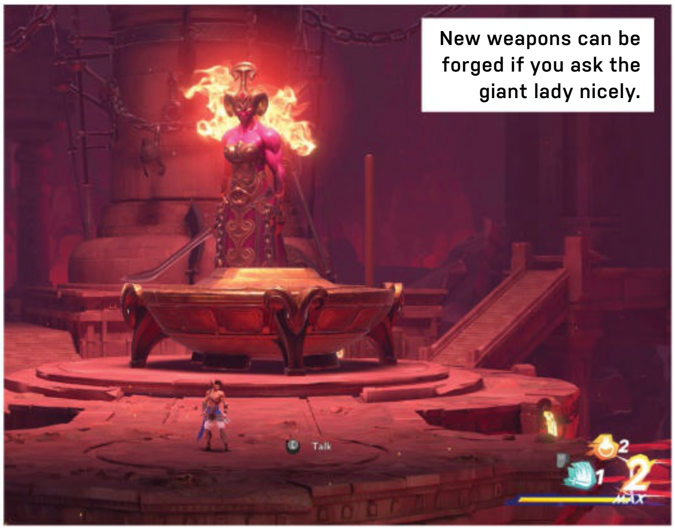

New weapons can be forged if you ask the giant lady nicely.

full torque, and the final 10 hours are like playing a completely different game.

At its best, then, *The Lost Crown*'s movement and platforming feels fantastic. It's responsive and fluid. Every jump, dash, and slide is quick and snappy, and controlling Sargon in trickier platforming segments never felt unwieldy.

The Lost Crown's combat is a highlight, too. Sargon can juggle enemies in the air, unleash a super move once his powers are fully charged or parry enemies to gain the advantage. It's a lot of fun, though in truth it's often easier to simply mash out the basic attack button and dodge with the occasional parry sprinkled in.

Despite all that, it's simply impossible to avoid a sense of frustration with how *The Lost Crown* pans out overall. The first half is painfully average, and while the gameplay improves, the story never recovers. If the whole of *The Lost Crown* played like its final few hours and had a more engaging narrative, Ubisoft would be onto a winner. As it is, it's at best a competent example of the art of metroidvania gameplay, and ultimately an opportunity missed. **–MOLLIE TAYLOR**

VERDICT
6

Prince of Persia: The Lost Crown

⊞ **PRINCELY** Response and fluid platform; Combat is full of fun.

⊟ **PLEBIAN** Takes forever to really get going; Flakey narrative.

RECOMMENDED SPECS CPU, Intel Core i7-6700, AMD Ryzen5 1600. GPU, NVIDIA GeForce GTX 960 or AMD Radeon RX 5500 XT. RAM, 8GB.

$50, www.ubisoft.com, rated 16

HEAD TO HEAD
WINNER

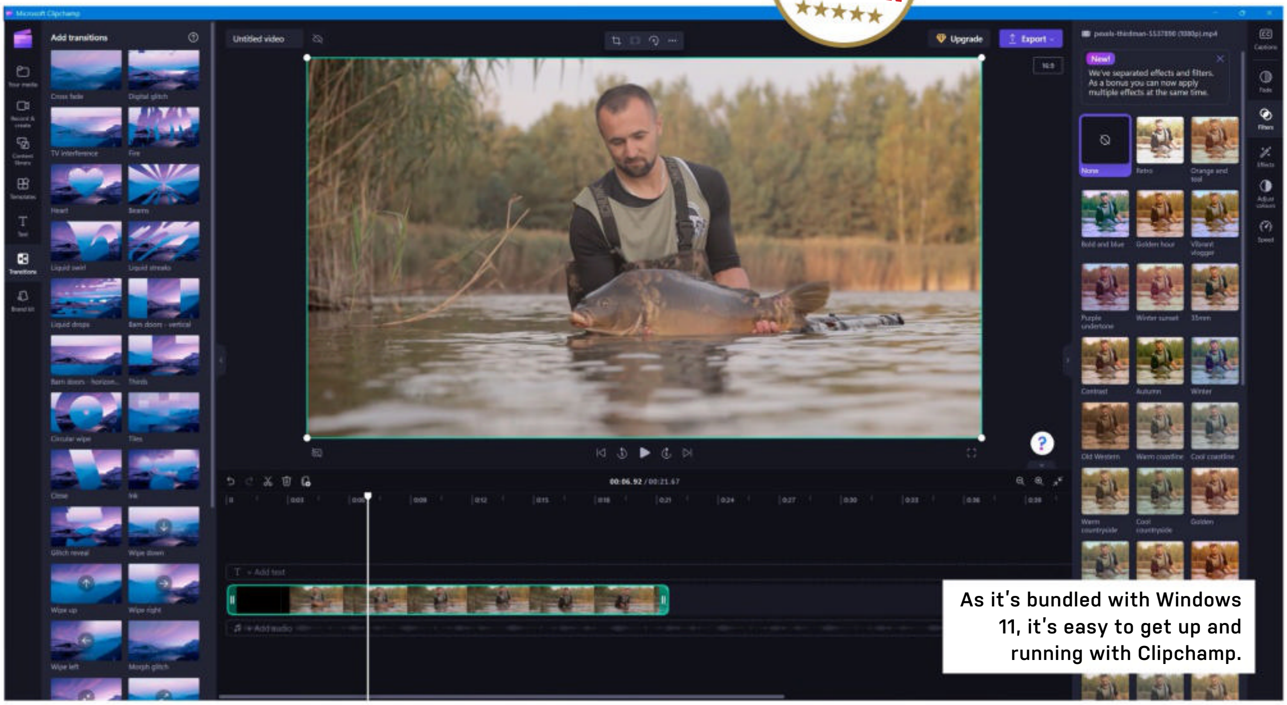
Microsoft Clipchamp
Add transitions
Untitled video
Upgrade
Export
New!
We've separated effects and filters. As a bonus you can now apply multiple effects at the same time.

As it's bundled with Windows 11, it's easy to get up and running with Clipchamp.

Microsoft Clipchamp vs Premiere Rush

There can be only one video editing winner

THERE'S PROBABLY a reason for the explosion in popularity of YouTube and the rise of multi-core PCs with GPUs happening at around the same time. Sure, you can shoot and edit video on a smartphone, but for a really professional-looking job, nothing beats getting your footage onto a fast SSD and viewing it on a big screen.

To make your edits, you're going to need software, and there's a lot available, from the market-leading Adobe Premiere Pro or Sony's Vegas Pro with their monthly subscriptions, to pay-once options, and free apps. It's the last category that has become especially interesting, as there are some capable apps available for absolutely nothing, such as the mighty DaVinci Resolve from Blackmagic design—the makers of some fearsome cinema cameras who use the software to bait users into their ecosystem—and apps like Shotcut and Lightworks that are good for beginners.

Then there are the more humble options; the perfect thing for quickly cutting together something to amuse your friends on Instagram. Windows 11 comes with such an option: the ambitiously named Microsoft Clipchamp, which has a premium version with a monthly subscription if you find it useful. Creative software powerhouse Adobe also gets in on the act with a few options, the Adobe Express free editor that runs in your browser, and Adobe Premiere Rush, which operates as a desktop app.

Both are good for getting straight to the business of editing your video clips, which you can drop straight into the apps from your PC. They offer a classic timeline-based editing interface, and can export your edited movie into a variety of formats. You can change the speed, apply color filters, and overlay effects such as transitions and titles. There's also the ability to separate the audio track from the video, and alter it or add music. Here's how these two titans of free video editing software stack up.

MICROSOFT CLIPCHAMP

The biggest benefit of Microsoft Clipchamp is that you probably already own it. The app was integrated into Windows 11 in 2022, but was developed by an Australian company back in 2014.

There are two versions of Clipchamp: one that runs in your browser, and a desktop app. The latter, available from the Windows Store if it's not already installed through an update, offers the standard cookie preference setting, so appears to be a web app rather than standalone, but it works offline once you've got it installed. The web version works in Edge and Chrome, but throws an error if you try to use it in Firefox—as Edge is Microsoft's own browser, and is based on the Chrome rendering engine, this makes sense, but is a little disappointing. At least it's not Edge exclusive, and is also available as an installable web app on a Chromebook.

Although Clipchamp is free, there are some premium features, such as 4K export, cloud backup, and additional visual and audio effects that you can unlock with a subscription of $11.99 a month. If your video project uses subscription features, your video will be watermarked on export. All the basic video editing tools are

© MICROSOFT

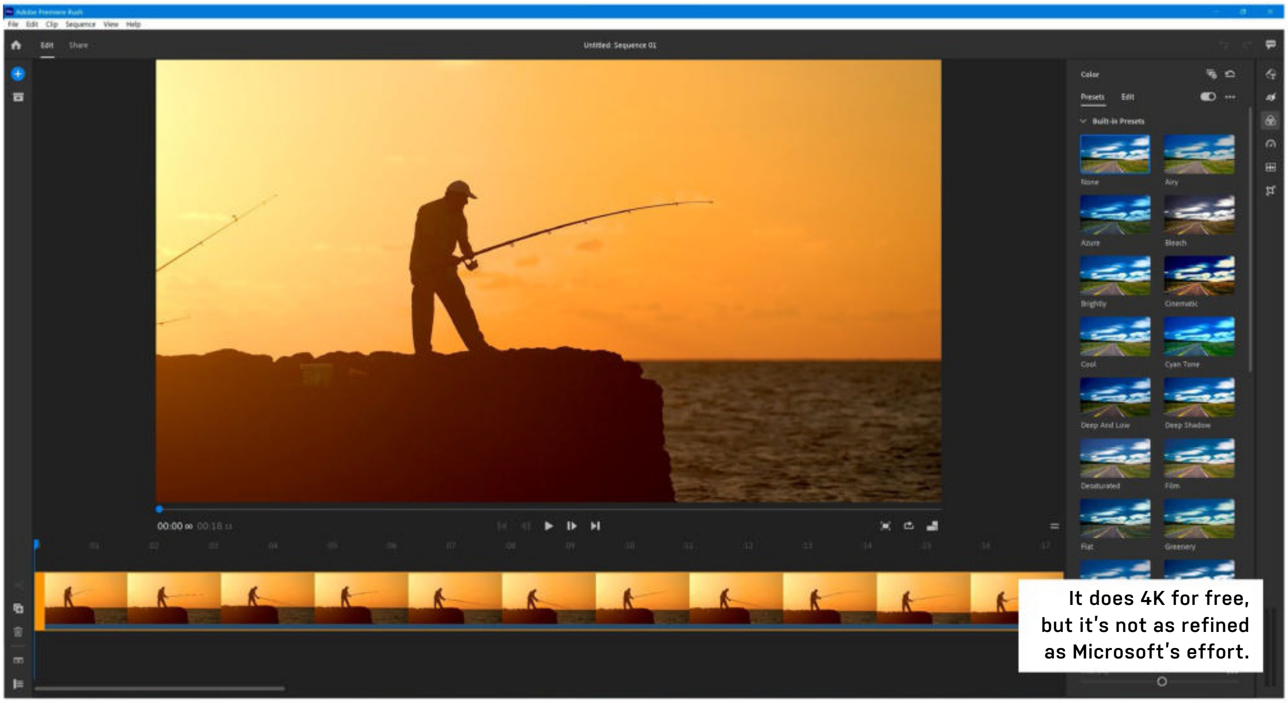

It does 4K for free, but it's not as refined as Microsoft's effort.

© ADOBE

included with the free version, however, and the timeline-based interface will be familiar if you've so much as looked at a video editing app. There are collapsible panels on either side of the preview window, which is an efficient use of space, as it allows a widescreen video clip to be as large as possible. The interface is touch-friendly, and there are keyboard shortcuts to zip between tools and views.

Alongside the standard method that sees you drag clips onto the timeline and dictate how they overlap and interact, there's an AI option that will create a movie from the raw materials. Once you've imported your clips into their tracks—you can point it to a folder or add them one by one with the mouse—you choose a style (or styles), a format, and a length of time, then a music style (there are royalty-free tracks), and let it do its thing. If you're the more hands-on type, or just don't like the result of the automatic edit, then you can use various templates, as well as recording directly from your webcam to add a picture-in-picture commentary to your video—useful if you're showcasing game footage or a software tutorial.

Clipchamp is a great place to start with video editing, especially as it's probably already on your PC. While its AI editing may not produce the results you'd expect from doing things by hand, it's faster, and you can always edit it later. You also don't have to use it—the templates and effects in the free version are a good choice if you want to learn about video editing.

ADOBE PREMIERE RUSH

Adobe has a lot of experience in making video editing apps, and alongside the subscription-based Premiere Pro, there's the cut-down Premiere Elements you can get for a one-off payment, and the Adobe Express free online video editor. Premiere Rush is an interesting addition, as it's a lightweight app that runs on mobile and desktop, and you can sync your work across devices via the Creative Cloud service, which means setting up an account via the CC Starter Plan, which also nets you 2GB of cloud storage.

Once you've installed the app, the interface is similar to Clipchamp and many other video editing apps. There's a large preview window, a timeline at the bottom, and bars to either side. Rush puts its media bin at the left, so you can drag clips onto the timeline easily, and all the graphics, transitions, and color presets on the right. One thing it doesn't do as well as Clipchamp is that the presets use an image of a road to show their effect rather than a grab of the clip you're working on. The Microsoft app has the better interface, both in the spread of its tools and the way its filters are displayed.

You can create your own presets too, or take manual control with a bank of sliders for exposure, color temperature, saturation, and more, and an 'Apply to All' button will copy your color tweaks to clips in your project. The actual editing interface has a maximum of seven tracks on its timeline: three for video and four for audio. The tracks automatically align and keep in sync, but it can be confusing that newly imported videos go straight into a primary track, and that extra video tracks are found above the primary one, while audio tracks are below. Once you've got that quirk set in your mind, you can drag between tracks, split, and duplicate happily. In contrast to Clipchamp, Rush can export a 4K video without watermarking or asking for a subscription.

While either app will do a great job of editing and exporting a 1080p video, Clipchamp has the edge in the interface department, so will be attractive to those making social media posts, but the 4K export capabilities of Rush will appeal to anyone with higher resolutions in mind. **–IAN EVENDEN**

VERDICT 8

Microsoft Clipchamp

+ **VIDEO** Bundled with Windows 11; Easy to use interface.

− **RADIO STAR** Advanced features locked behind a subscription.

Free/$11.99 a month

VERDICT 7

Adobe Premiere Rush

+ **RUSH** 4K without a sub, plenty of presets and manual options.

− **HUSH** Interface doesn't have the refinements of Microsoft's app.

Free

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> Water Bored
> Replace Tablet Battery
> Building Ideas
> Super Bad?

Trickle Down Effect

I may be wrong, but I don't recall seeing a feature on water cooling recently—real water cooling with a custom loop. It's something that I feel is just outside of my reach, and even though I could YouTube and Google it, I think having a several-page feature that compares the options and pathways to completing it might be nice.

Thanks for producing the magazine. I'm an old enthusiast that has enjoyed the title since the days of Gordon. If any of you speak to him, wish him well from me. I've been very sad to see him recently, but I'm glad that he is continuing his fight. **–S. Rakestraw**

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: Mr Rakestraw, you are correct—we haven't done a lot on water cooling since our 'Water cooling Made Easy' cover feature in June 2022. But as we now have an expert in the field, Zak Storey, working for us again, I'm going to bring this oversight up with him in the hope that he's able to get us up to date.

I went to our head office as part of that feature in order to get a masterclass from Zak at the time, who was using one of Corsair's liquid cooling kits to show us how to fill reservoirs, bend tubing, and add water blocks. It was a fascinating experience, and one I'd encourage any PC builder to try, even if the idea of mixing liquid and PC components can seem daunting. It's easier than it looks, and while bending pipes can take some getting used to at first, it's actually pretty hard to get anything majorly wrong when using the latest kits.

The attraction of liquid-cooled PCs is that you can create a virtually silent system, which also looks cool, and can be overclocked like crazy. With Intel's 14th-gen chips and the new PCIe 5.0 drives running particularly hot, these seem to be ideal candidates for the best cooling possible, so the time seems right for us to do something on this. I'll talk to Zak and report back.

As for Gordon Ung, who is previously of this parish, those on Reddit or following PCWorld may have seen that he announced his cancer diagnosis a few months ago. You can follow him on social media @GordonUng, and also see him talking about it on PCWorld's YouTube channel (tinyurl.com/n5h78x4y). Everyone at *Maximum PC* wishes him all the best, and hopes his treatment is speedy and successful.

Surface Attraction

I have a powerful desktop with a double wide display that I do not intend to replace. I also have an old Microsoft Surface. I've been planning to use the Surface as a second machine to do White Boarding with a stylus when I'm teaching (as I'm a professor). Sadly, though, the Surface is coming to the end of its battery life now. Obviously, I can use it plugged in at the desk, just not off the battery.

But is there anything realistic to be done to save it as a portable? Or am I looking at replacing and recycling it? **–R. Matlow**

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: You've not mentioned which version of Microsoft's Surface you're using, but the good news is you can do a battery replacement on most models if you have the time and expertise. Our good friends at iFixit have multiple guides on how to do this (tinyurl.com/ez65fw6m), and they're rated between 'moderate' to 'difficult', depending on which model you're looking at. You can buy the tools you'll need from iFixit directly, while replacement batteries are pretty easy (and cheap, at under $30) to find on eBay.

If you do decide to ditch the Surface, or better yet, donate it (see computerswithcauses.org) or give it to your school or college, I would recommend looking at the Samsung Galaxy Tab S9 as a replacement. Yes, it's Android rather than Windows-based, but I've used this tablet for the last six months or so, and love its ability to work with the included stylus, integrate well with

↘ submit your questions to: editor@maximumpc.com

Windows, and turn into a basic laptop replacement with Samsung's DeX mode. It also boasts a beautiful OLED display, which is likely to put your Surface in the shade. It's included on all models in the range for the first time.

Art vs Science

I very much enjoy *Maximum PC*, and I submitted a two-year renewal several weeks ago. I've been looking to build a living room gaming PC, so I looked back at the August 2022 edition, which featured such a machine on page 16. Thank you for giving me some inspiration—maybe it's time to do an updated build guide for 2024 though?

In the meantime, there are some things I'd like to see in future Build It articles: BIOS/UEFI advanced settings; Using hard line custom tooling, like tools, measuring, bending, adapters, etc; RGB (I know you don't like it) manufacturer control programs; and how to synchronize the display with free or inexpensive programs, such as Signal, Jack's RGB, etc.

As for the rest of the magazine, I like the articles that I learn computer-related things from. I don't care about your extra-length 'science' articles about supercomputers, photorealism, etc. Just keep it to the PCs, please!

Thank you for letting me vent. **–D. King**

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: Thanks for your feedback on our 'Ultimate Couch Gaming' issue, the build of which I put together using NZXT's very Xbox Series X-looking H1 v2 case. You're right, we're probably due a refreshed living room PC build, as that was done before the current generation of GPUs arrived later that year, and it would be interesting to see which small cases would work best, given how large they are and how much air they draw in for cooling.

In terms of your Build feature suggestions, I'll take these one by one. We often run through BIOS settings in the final part of our builds, pointing out which settings you want to focus on and toggle once you get power to your machine. As covered in the first letter, we're overdue a feature on water cooling, of which piping will be a part of it, so stay tuned for that later this year. As for RGB, yes, we've had other readers asking about how to sync up all their lighting (and not just in their PC—their entire office or gaming room as well), so we need to write something on this. While our cover features are often adorned with gorgeously synched lighting, in my personal PC, I often put up with different components being out of sync because I can't be bothered to fiddle around with different apps on my computer to get them connected. But never let it be said that we on *Maximum PC* don't like a bit of RGB lighting!

In terms of our features, thanks for the feedback. It's really useful to get readers' thoughts. Obviously, we try to balance the magazine so that there's a good mix of content in there for all of our readers, but I certainly feel like I've been reducing the number of forward-looking 'science' pieces, as you describe them, in favor of more practical PC-based articles. If you're missing the former, though, please write in and tell me!

The launch of Nvidia's SUPER cards has shaken up the GPU market.

RTX 4070 SUPER

© NVIDIA

Super Fatigue

I've been holding off on buying a graphics card for the whole of last year, given the sky-high prices of AMD and particularly Nvidia's GPUs, so I was eagerly awaiting the rumored announcement of Nvidia's SUPER cards. I wish that there was one of them available for under $500, as my mind will never be able to cope with the idea of a graphics card costing more than half a thousand dollars, but I'm ready to bite the bullet with the 4070 SUPER at $600, given that I've been waiting so long. But my question to you: is that the best option, or should I wait a bit longer? **–J. Evans**

EDITOR-IN-CHIEF, GUY COCKER, RESPONDS: The danger with waiting to buy anything tech-related is that it's an always-moving target, and if you don't take the plunge, you'll just end up holding off forever. That said, now is a better time than any since the launch of the current generation of GPUs at the end of 2022 to be buying a card.

There is a card available at close to your chosen price, and that's the RTX 4070, which has dropped following the announcement of the SUPER edition. We haven't had the SUPER edition in to compare the two, but on paper it has 20 percent more cores than the RTX 4070, so it *should* be worth the extra $60-70 that it costs at time of writing, if you can stretch your budget that far.

Another option is to take advantage of the competition Nvidia is bringing to the space and the resulting AMD price drops. The RX 7700 XT can now be found for $399 (about $50 less than launch), while the RX 7800 XT is holding firm at $499—both are well within your budget. It doesn't sound like you're willing to go up to $800, but the top-end RX 7900 XTX dropped to that around Christmas time. They don't have the fancy DLSS 3.5 bells and whistles that Nvidia's cards have, but if all you care about is raw performance, then AMD will serve you well.

Good luck, and let us know how you get on!

THE BUILDS

THIS MONTH'S STREET PRICES...

THIS ISSUE'S been an interesting one for blueprints. What we're seeing across the industry is a slow but steady increase in SSD prices across the board. Almost every single SSD pick we had from last issue has increased by $20-30. This is likely down to those bumper New Year and Black Friday deals finally subsiding. However, we're mildly concerned that this may be a continuing trend across the industry, depending on how NAND pricing continues moving forward. It's one to watch, for sure.

Our budget builds for this issue have remained reasonable, all things considered. Changes were minimal—we've swapped a few components out across the two builds, both of those cases being the memory across the systems and the primary SSD. We've opted for the 1TB Crucial P3 Plus PCIe 4.0 drive for our OS, as the MSI Spatium M450 had increased significantly.

We also changed up the memory. Teamgroup's Elite Plus DDR5 @ 4800 had climbed by $20 for AMD, so we swapped it for the non-heatsinked Elite variant for a slightly lower price. The Patriot Viper Venom @ 5200 kit seemingly disappeared from stock entirely, so we've gone for the 5600 kit for only $3 more.

Intel's build also saw price cuts on its motherboard, graphics card, and processor, pulling the overall price down by around $35. The same can't be said for AMD, however, with increases bumping the price up by $20, with a number of AMD's mainstream processors rising in price compared to last issue. This could be a reaction to Intel's 14th gen chips not being as popular as first thought, and AMD responding in kind. While that is speculative, it is interesting to note that there was no AMD keynote at this year's CES. If you can hold off, we'd recommend waiting for Computex in June. It's likely that both AMD and Intel will be launching new chips, which will either bring these prices down or give us new CPUs with better performance at similar prices.

AMD INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$95
PSU	600W Thermaltake Toughpower GX2 80+ Gold	$65
Mobo	Asus Prime A620-Plus WiFi ATX AM5	$140
CPU	AMD Ryzen 5 7600	$199
GPU	XFX Speedster Swift 210 Core RX 7600 8GB	$270
RAM	16GB (2x8GB) TEAMGROUP Elite @ 4800 **NEW**	$46
SSD 1	1TB Crucial P3 Plus PCIe 4.0 M.2 **NEW**	$58
SSD 2	1TB Teamgroup T-Force Cardea Z44L PCIe 4.0 M.2	$56
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $961

INTEL INGREDIENTS

PART		PRICE
Case	Corsair 4000D Airflow	$95
PSU	600W Thermaltake Toughpower GX2 80+ Gold	$65
Mobo	ASUS Prime B760-Plus ATX	$100
CPU	Intel Core i5-12400F	$144
GPU	Gigabyte Windforce OC RTX 4060 8GB	$300
RAM	16GB (2x8GB) Patriot Viper Venom DDR5 @ 5600 **NEW**	$65
SSD 1	1TB Crucial P3 Plus PCIe 4.0 M.2 **NEW**	$58
SSD 2	1TB Teamgroup T-Force Cardea Z44L PCIe 4.0 M.2	$56
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32

Approximate Price: $915

MID-RANGE

IT'S A SIMILAR STORY in our mid-range build, with AMD's overall price going up by around $30, and Intel decreasing by $20. Again, this is speculative, but that reaction to the 14th generation of Intel chips has been fairly muted from the community and tech media alike. That said, similar to our budget builds, there's a few extra changes across both systems this issue, with the focus being SSDs again. Sadly, the Lexar NM790 and Gigabyte Aorus Gen5 from last issue have taken a hit. In fact, it's got so bad that we've had to drop down the capacity on our secondary drive, back to 1TB from the 2TB drives we were running last issue.

The 1TB Corsair MP600 PRO LPX is our primary pick for the OS drive. Clocking in at $80, it's an impressive little number. Coming with a low-profile heatsink (initially designed for the PS5), it's a solid pick for any gaming PC, and its performance is top tier, matching the beefier MP600 Pro that launched earlier in Corsair's MP series life cycle. It clocks in sequential reads at 7,100 MB/s and writes at 5,800 MB/s. To back that up, we've gone with the lower-specced Adata Legend 800 PCIe 4.0 SSD for our backup drive. At just $58, it's an impressively valued proposition, and will be more than enough for all of our gaming needs.

One additional change comes in with our AMD build, and that's in the form of the graphics card. PowerColor's Fighter OC Radeon RX 7700 XT increased in price by a staggering $60, which is a bit too much for us to stomach, so we've opted to ditch it in favor of the ASRock Challenger OC variant. You can expect similar performance from the two cards (the difference seems to be within 1-2 percent, depending on the title), and 1440p gaming is in the high 80s comfortably with it (particularly if you enable FSR). It's slightly more expensive than last issue's card, coming in at $10 more, but—and it's a big but—this is currently the cheapest 7700 XT that money can buy.

AMD INGREDIENTS

PART			PRICE
Case	NZXT H7 Flow		$118
PSU	750W Corsair RM750e - 80+ Gold		$100
Mobo	MSI PRO X670-P WiFi ATX - AM5		$221
CPU	AMD Ryzen 7 7700X		$349
Cooler	Be Quiet! Pure Loop 2 - 240mm AIO		$95
GPU	ASRock Radeon RX 7700 XT Challenger OC	NEW	$430
RAM	32GB (2x16GB) Corsair Vengeance RGB DDR5 @ 6000		$110
SSD 1	1TB Corsair MP600 PRO LPX M.2 PCIe 4.0 SSD	NEW	$80
SSD 2	1TB Adata Legend 800 M.2 PCIe 4.0 SSD	NEW	$58
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)		$32
Approximate Price:			**$1,593**

INTEL INGREDIENTS

PART			PRICE
Case	NZXT H7 Flow		$124
PSU	750W Corsair RM750e - 80+ Gold		$100
Mobo	MSI Z790 Gaming Pro WiFi ATX		$210
CPU	Intel Core i5-14600KF		$303
Cooler	EK AIO Basic 360 - 360mm AIO		$129
GPU	MSI Ventus 3X OC RTX 4060 Ti 8GB		$405
RAM	32GB (2x16GB) Teamgroup T-Force Vulcan DDR5 @ 6000		$87
SSD 1	1TB Corsair MP600 PRO LPX M.2 PCIe 4.0 SSD	NEW	$80
SSD 2	1TB Adata Legend 800 M.2 PCIe 4.0 SSD	NEW	$58
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)		$32
Approximate Price:			**$1,528**

THIS MONTH'S TURBO BUILDS are a tale of price increases and decreases. Take our case pick. Up by $35. Then the PSU? Down by $60, motherboard for AMD? Up by $30. CPU (before we changed it) up by $80. AMD RAM? Up by $40. Yep, it's been a doozie of a build list, that's for sure. Still, we've managed to throw in a few clever product tweaks to keep the price consistent, meaning our AMD build has decreased by $10, without sacrificing performance, and our Intel build has shot up by $45 compared to last issue (which at high end prices, is fairly reasonable).

The biggest change comes in the form of our CPU pick for the AMD build. We've ditched the 7950X3D chip and opted for its more traditional, non-3D cousin, the 7950X. This is still a potent processor, particularly if you're looking at any form of multi-threaded applications or processing. Anything rendering based, and this absolutely rips through it. Although it's not as potent at gaming as either the 7950X3D or Intel's offerings, it's well within five to ten percent of its 3XD counterpart for fps in-game at 1080p, and 10-15 percent to Intel's top dog 14900K, although this is dependent on memory frequency (we recommend pairing it with a 6000 MT/s kit to get the most out of that infinity fabric).

Speaking of memory, for AMD we've done exactly that, ditching the Corsair Vengeance kit (which had skyrocketed by $60) in favor of the Teamgroup T-Create Expert 6000 C34 kit. It's not as impressive as last month's kit, as it's got a C34 latency rather than a C30, but given the closest Vengeance kit price-wise came in at C40, it was the smarter pick of the two at 6000 MT/s. It's also available in black or white at the same price, and even compared to last issue comes in at $60 cheaper. No RGB, though.

We kept the Intel build solidly positioned, with only one major change. Interestingly, our left-field 48GB memory pick has stayed steady on the pricing (compared to pretty much every other build and SKU), so we kept it in. The change? The GPU. PNY's Verto RTX 4080 shot up by $40 again, so we've swapped that out for MSI RTX 4080 Ventus 3X, which is the better card anyway, and comes in at exactly the same price as last issue.

AMD INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$180
PSU	Be quiet! Straight Power 12 1200W - 80+ Platinum	$180
Mobo	Asus Prime X670E-Pro WiFi ATX - AM5	$300
CPU	AMD Ryzen 9 7950X **NEW**	$549
Cooler	NZXT Kraken Elite 360 RGB - 360mm AIO	$287
GPU	Sapphire Nitro+ RX 7900 XTX 24GB	$1,080
RAM	64GB (2x32GB) Teamgroup T-Create Expert @ 6000 C34 **NEW**	$180
SSD 1	2TB Corsair MP700 PCIe 5.0 M.2	$260
SSD 2	4TB Teamgroup T-Force Cardea Z44Q PCIe 4.0 M.2	$230
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$3,278**

INTEL INGREDIENTS

PART		PRICE
Case	Phanteks Enthoo Pro 2 Tempered Glass	$180
PSU	Be quiet! Straight Power 12 1200W - 80+ Platinum	$180
Mobo	Gigabyte Z790 Aorus Elite AX-W ATX	$362
CPU	Intel Core i9-14900KF	$544
Cooler	NZXT Kraken Elite 360 RGB - 360mm AIO	$287
GPU	MSI Ventus 3X OC RTX 4080 16GB **NEW**	$1,150
RAM	48GB (2x 24GB) G.Skill Trident Z5 RGB DDR5 @ 6800 CL34	$170
SSD 1	2TB Corsair MP700 PCIe 5.0 M.2	$260
SSD 2	4TB Teamgroup T-Force Cardea Z44Q PCIe 4.0 M.2	$230
OS	Windows 10 Home 64-bit OEM (Windows 11 Compatible)	$32
Approximate Price:		**$3,395**

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Holiday issue following the December issue, by Future US LLC, 130 West 42nd Street, 7th Floor, New York, NY 10036. USA. Website: www.futureus.com. Future US LLC also publishes *Mac|Life*, and *PC Gamer*. Entire contents copyright 2023, Future US LLC. All rights reserved. Reproduction in whole or in part is prohibited. Future US LLC. is not affiliated with the companies or products covered in *Maximum PC*. Reproduction on the Internet of the articles and pictures in this magazine is illegal without the prior written consent of *Maximum PC*. Products named in the pages of *Maximum PC* are trademarks of their respective companies.

PRODUCED IN THE UNITED STATES OF AMERICA.

UNEARTH THE FULL STORY BEHIND THE LEGEND OF ZELDA

Explore every game in the series and what makes each one so great. If you're a fan, you'll enjoy the trip through time. If you're a newcomer, we offer you this sage advice: "It's dangerous to go alone, take this!"

OLD!
retro GAMER
PRESENTS
100% UNOFFICIAL
THE STORY OF
ZELDA
DISCOVER THE COMPLETE HISTORY OF HYRU
BRAND NEW
TEARS OF THE KINGDOM REVEALED
EVERY GAME IN THE SERIES EXPLORED
TIMELINES EXPLAINED
BEHIND-THE-SCENES ACCESS
Digital Edition

ON SALE NOW

FUTURE

Shop easily online at:

magazinesdirect.com

bit.ly/3xbOGXq

DIGITALSTORM

PC GAMER
EDITOR'S CHOICE
MAXIMUM PC
KICK ASS!

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM
Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside

The Digital Storm Logo, and 'World's Most Advanced PCs' are trademarks of Digital Storm.Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.